Investigations into the effects of neuromodulations on the BOLD-fMRI signal

Melissa May Maczka

Department of Statistics

University of Oxford

A thesis submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Michaelmas 2013

Supervisors: Jonathan Marchini, Mark Woolrich and Chris Martin

Acknowledgements

This thesis was made possible by my supervisors, Dr. Chris Martin, Dr. Mark Woolrich and Dr. Jonathan Marchini. With their guidance, I have become a scientific researcher able to formulate hypotheses, test hypotheses by running experiments *in silico* or *in vivo*, and then critically evaluate experimental evidence. The majority of this learning has taken place in the Department of Statistics, which provided me with a quiet and comfortable working space for three years.

The aspect of my DPhil that has proven most dependent on the kindness of others is my training to become an *in vivo* experimentalist. This training, which was conducted at the Radiobiology Research Institute, would not have been possible without the help of a number of their academics and support staff, most notably Dr. Nicola Sibson, Dr. Aisling Spain, Dr. Sebastian Serres and Dr. Yvonne Couch.

Much of this thesis relies on experimental data collected by collaborators. Dr. Maria Franceschini and Dr. Harsha Radhakrishan, both of Athinoula A. Martinos Center for Biomedical Imaging, Harvard University, very kindly provided the anaesthetics dataset analysed in Chapters 2, 3 and 4 in this thesis. The LPS dataset analysed in Chapter 6 was provided by Prof Bernhard Rosengarten of Universitätsklinikum Gieund Marburg.

Other individuals who have contributed to my thesis writing include my Transfer of Status, Confirmation of Status and Final Examination examiners Prof Ying Zheng, Dr. Eugene Duff, Dr. Clare Howarth and Dr. Adrian Groves. Their feedback was instrumental in helping me shape my thesis.

Finally, I feel certain that I was only able to produce the work in this thesis because of the research skills and knowledge I gained at the Doctoral Training Centre in 2009-10. I am also thankful to the DTC for awarding me with the EPSRC Studentship that has allowed me to pursue postgraduate training for the past four years.

Investigations into the effects of neuromodulations on the BOLD-fMRI signal

Abstract

Melissa May Maczka, St. Cross College
Submitted in partial fulfillment of the requirements for the
degree of Doctor of Philosophy, Michaelmas 2013

The blood oxygen level dependent functional MRI (BOLD-fMRI) signal is an indirect measure of the neuronal activity that most BOLD studies are interested in. This thesis uses generative embedding algorithms to investigate some of the challenges and opportunities that this presents for BOLD imaging.

It is standard practice to analyse BOLD signals using general linear models (GLMs) that assume fixed neurovascular coupling. However, this assumption may cause false positive or negative neural activations to be detected if the biological manifestations of brain diseases, disorders and pharmaceutical drugs (termed "neuromodulations") alter this coupling. Generative embedding can help overcome this problem by identifying when a neuromodulation confounds the standard GLM. When applied to anaesthetic neuromodulations found in preclinical imaging data, Fentanyl has the smallest confounding effect and Pentobarbital has the largest, causing extremely significant neural activations to go undetected. Half of the anaesthetics tested caused overestimation of the neuronal activity but the other half caused underestimation. The variability in biological action between anaesthetic modulations in identical brain regions of genetically similar animals highlights the complexity required to comprehensively account for factors confounding neurovascular coupling in GLMs generally. Generative embedding has the potential to augment established algorithms used to compensate for these variations in GLMs without complicating the standard (ANOVA) way of reporting BOLD results.

Neuromodulation of neurovascular coupling can also present opportunities, such as improved diagnosis, monitoring and understanding of brain diseases accompanied by neurovascular uncoupling. Information theory is used to show that the discriminabilities of neurodegenerative-diseased and healthy generative posterior parameter spaces make generative embedding a viable tool for these commercial applications, boasting sensitivity to neurovascular coupling nonlinearities and biological interpretability. The value of hybrid neuroimaging systems over separate neuroimaging technologies is found to be greatest for early-stage neurodegenerative disease.

Contents

Contents iii

1 Introduction 1

1.1 Traditional approaches for assessing the effects of neuromodulations on BOLD-fMRI . . . 3

1.2 Thesis Summary: Using generative embedding to assess the effects of neuromodulations on BOLD-fMRI signals . . . 9

1.2.1 Neuromodulations as confounding factors for BOLD-fMRI studies . . . 10

1.2.1.1 Chapter 2 . . . 10

1.2.1.2 Chapter 3 . . . 12

1.2.1.3 Chapter 4 . . . 13

1.2.2 Neuromodulations as tools for neuroimaging investigations . . . 13

1.2.2.1 Chapter 5 . . . 13

1.2.2.2 Chapter 6 . . . 14

1.2.2.3 Chapter 7 . . . 15

1.3 Novel contributions to the field . . . 16

2 Testing BOLD-fMRI general linear model sensitivity to anaesthetic neuromodulation 18

2.1 Anaesthetics dataset . . . 19

2.1.1 Experimental protocol . . . 21

2.2 GLM analysis . . . 24
2.3 Methods . . . 26
2.4 Results and Conclusions . . . 28

3 Testing BOLD-fMRI GLM sensitivity when accounting for anaesthetic neuromodulation of neurovascular coupling 40

3.1 Generative modelling . . . 41
3.1.1 Single-modality DCMs . . . 41
3.1.2 Combining DCM-EEG and mDCM-fMRI . . . 47
3.1.2.1 Details of Combined DCM . . . 50
3.1.2.2 Modifications to Sotero et al. linking model . . . 55
3.1.2.3 Combined DCM parameters . . . 55
3.2 Methods . . . 59
3.2.1 Combined DCM fitting . . . 59
3.2.2 Simulating [ΔHbR] data . . . 65
3.2.3 Embedding and Inference . . . 67
3.3 Results . . . 67
3.4 Conclusions . . . 72
3.4.1 Capturing nonlinearities in GLM . . . 73
3.4.2 Tailoring HRFs in GLM . . . 75
3.4.3 Correcting for neuromodulations using generative embedding . 76

4 Insights into the effects of anaesthetic neuromodulation on neurovascular coupling 78

4.1 Qualitative observations . . . 79
4.2 Demonstrating how anaesthetics modulate hemodynamic linearity and impulse response shape . . . 87
4.2.1 Methods . . . 89
4.2.2 Results . . . 90

4.2.3 Discussion . . . 91
4.3 Understanding which specific biological actions of anaesthetics confound GLM analyses . . . 95
4.3.1 Methods . . . 95
4.3.2 Results and discussion . . . 97
4.4 Overall discussion . . . 98

5 Do neurodegenerative diseases alter neurovascular coupling? 103
5.1 BOLD-fMRI as a tool for diagnosing/monitoring/developing drugs for brain diseases . . . 104
5.1.1 Neurodegenerative diseases . . . 106
5.1.2 Predictions . . . 109
5.2 Methods . . . 113
5.2.1 Experimental data collection . . . 113
5.2.1.1 Experimental protocol . . . 113
5.2.2 Simple analysis: linear regression of shapes features . . . 115
5.3 Results . . . 116
5.4 Discussion . . . 117

6 Can we exploit neuromodulatory effects to better diagnose neurodegenerative disease? 120
6.1 Description of neurodegenerative disease dataset . . . 121
6.2 Generative Embedding: Information-theoretic statistics . . . 123
6.2.1 Information Theory . . . 124
6.2.1.1 Synergy . . . 125
6.2.2 Bias in information-theoretic calculations . . . 126
6.2.2.1 Choosing an algorithm to estimate densities with low bias . . . 126
6.2.2.2 k-nearest neighbours . . . 129

6.3 Methods . . . 132
6.3.1 Using k-nearest neighbours to estimate densities including categorical variables . . . 134
6.3.2 Finding optimal values for k_{fixed} . . . 137
6.4 Results . . . 138
6.5 Discussion . . . 139
6.5.1 MI . . . 143
6.5.2 Synergy . . . 149
6.5.3 Technological viability . . . 150

7 Commercialising generative embedding for investigating neurovascular coupling 153
7.1 Introduction . . . 153
7.2 Potential commercial products . . . 154
7.2.1 Imaging modalities . . . 156
7.3 Markets . . . 158
7.3.1 Imaging in the clinic today . . . 159
7.3.2 Imaging in drug development today . . . 161
7.3.3 Market requirements . . . 163
7.4 Recommendations . . . 165

8 Conclusions 168
8.1 Evaluation of Chapters 2 - 4 . . . 171
8.1.1 Future work . . . 175
8.2 Evaluation of Chapters 5 - 6 . . . 176
8.3 Evaluation of Chapter 7 . . . 177
8.3.1 Neurovascular imaging modalities of the future . . . 178
8.4 Summary . . . 179

A BOLD GLM 181

B Fitting DCMs **184**

B.1 State variable integration 184

B.2 Variational Bayes algorithm 185

C Review of Linking Models **188**

C.1 Buxton et al. and Babajani et al.'s models 189

C.2 Sotero et al.'s model 190

D Literature review: Bias correction schemes **192**

E Future modifications to Combined DCM **195**

E.1 Separating presynaptic excitability and inhibition 195

E.2 Capturing neural adaptation 196

E.3 Capturing the post-stimulus undershoot more correctly 200

E.4 Reducing model complexity 202

References **204**

Chapter 1

Introduction

The aim of many psychology and cognitive neuroscience studies is to determine the differences in brain activities between groups of subjects differentiated by a factor, in this thesis referred to as an interest neuromodulation. These interest neuromodulations may either be endogenous, i.e. arising from physiology (pathological or otherwise), or they may be exogenous, i.e caused by administration of a substance, such as a pharmacological drug.

Blood oxygen level dependent (BOLD) functional magnetic resonance imaging (fMRI) is the neuroimaging modality of choice for many researchers when conducting these types of psychology and cognitive neuroscience studies because it is non-invasive and boasts high spatial resolution. However, because the BOLD signal measures neuronal activity through a hemodynamic veil (i.e. through its effects on blood dynamics, see Box 1), which is known to be susceptible to neuropharmacological effects, disease and disorder, it is difficult to establish which differences in the BOLD signals between the experimental groups result from differences in the underlying neuronal activities (which we seek to establish), and which result from differences in neurovascular coupling (a secondary and confounding effect of the neuromodulations, see Box 2). If the differences in the BOLD signal resulting from differences in neurovascular coupling are not accounted for correctly, the differences in the neuronal activities between the

experimental groups will almost certainly be inferred incorrectly. For example, patients with chronic cerebral ischemia may have sustained dilation of blood vessels distal to those that are occluded, meaning that brain tissues near the occluded area are sufficiently perfused to function normally despite experiencing smaller increases in hemodynamics between resting and active states than control subjects [40][42]. An fMRI study comparing the neuronal activities of these two experimental groups might conclude that the subjects with chronic cerebral ischemia have impaired neuronal activity in that brain area, when in fact changes to neurovascular coupling are maintaining their neuronal activities at the same level as for the control group. To avoid errors like this, it is necessary to develop a method for predicting neuronal activity from the BOLD signal that is able to account for these sorts of changes in neurovascular coupling.

General linear modelling (GLM) forms the basis of the algorithms used to identify regions of neural activation and deactivation in most human BOLD studies today. It involves creating a function that predicts temporal BOLD signal using knowledge of the stimulation applied and a set of basis functions, and then identifying brain regions where this prediction has highly positive (activation) or negative (deactivation) correlation with observation. 62% of journal papers randomly drawn from the BOLD-fMRI literature use the simplest GLM, which is inflexible with regard to neurovascular coupling, and the majority of the remaining journal papers use variants of this simple GLM that have limited success in regressing out differences in neurovascular coupling [72]. For these reasons, there is considerable concern within the fMRI community that the conclusions of many published fMRI studies are incorrect. What is more, it is frequently reported that hemodynamic responses to neuronal activity vary across subjects and even brain regions [195][4], meaning that neuromodulations are not the only cause for BOLD studies to be confounded. This highlights a general problem for BOLD imaging: a large number of fMRI studies continue to generate novel information about the

complex functional organisation of the brain, yet considerably less attention is being invested in validating and improving all aspects of its use. For example, by 2008 only 5% of journal papers that included the keyword "fMRI" or similar were concerned with the basic science behind the fMRI signal rather than using it to probe a psychological or cognitive neuroscience question [103].

1.1 Traditional approaches for assessing the effects of neuromodulations on BOLD-fMRI

As already mentioned, the aim of most BOLD-fMRI experiments is to determine the differences in neuronal activities between control and neuromodulated groups. However, neuromodulations may confound these experiments if they alter neurovascular coupling because the GLM may attribute these differences in neurovascular coupling to the interest variable (differences in neuronal activity caused by the interest neuromodulation, see Fig.1.3), leading to either a Type 1 or Type 2 error:

1. A Type 1 error (false positive) will occur when there are no significant differences in brain activations between the two groups, but the effect of the neuromodulation on neurovascular coupling makes it appear as if there are significant differences.
2. A Type 2 error (false negative) will occur when there are significant differences in brain activations between the two groups, but the effect of the neuromodulation on neurovascular coupling masks them from the GLM analysis.

The traditional approach for investigating aspects of neurovascular coupling of relevance to BOLD GLM has been to extract shape features (usually response amplitude or power) from averaged recordings of neuronal activity

Box 1: The physics of BOLD-fMRI

In 1990, Ogawa et al. showed that the deoxyhemoglobin (HbR) in blood can act as an endogenous contrast agent for neural activation in MRI because [133]:

- It is paramagnetic and so causes magnetic field inhomogeneities that dephase water proton spins (see Fig.1.1). This dephasing results in a reduction in MR signal intensity.
- The HbR concentration in an area of elevated neural activity typically drops, owing to a large volume of freshly oxygenated blood arriving to meet increased metabolic demands. This is made possible by dilation of local blood vessels resulting from a cascade of signals that relax the muscles in the arteriole walls in response to neural activation (see Box 2) [186].

In contrast, oxyhemoglobin (HbO) is diamagnetic and so does not alter magnetic susceptibility. Therefore, the BOLD signal reflects the flux of deoxygenated blood in a brain area, which is an indirect marker of its neural activity.

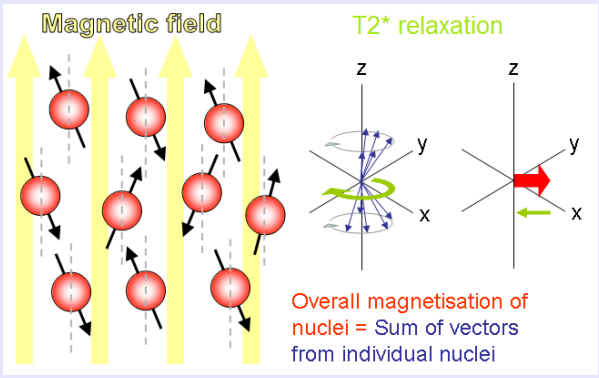

Figure 1.1: Explanation of BOLD imaging. Left: When water protons are exposed to a main magnetic field in the z-plane, their magnetisation either aligns roughly parallel or anti-parallel to it. When they are excited by a pulse of radio-frequency, their magnetisation is rotated into the x-y plane, with each proton rotating in phase about the z-axis. The magnetisation in the x-y plane gives rise to alternating voltage that is detected by a receiver coil. Centre: In the presence of HbR, the main magnetic field becomes inhomogeneous, causing dephasing of the water protons, known as T2* relaxation (blue arrows). Right: This results in a reduction in their overall magnetisation (from the red arrow, to the sum of the red and green arrows), and therefore reduction in MR signal. Figure reproduced from [31].

Box 2: Neurovascular coupling

Although the brain accounts for 20% of a human's energy needs [158] and consumes 25% of a human's oxygen intake [107], there is little storage available for these nutrients in the brain. Therefore, to ensure that sufficient amounts of these nutrients are available to support neuronal activity[a], they need to be delivered to active brain areas quickly. This delivery, called functional hyperemia or neurovascular coupling, is achieved by modulation of cerebral blood flow (CBF) via dilation/constriction of the smooth muscle cells lining the local arterial vessels (as this alters the resistance to CBF)[b] [151]. The traditional view is that there is a causal link between neural energy demand and CBF, perhaps via O_2, CO_2 or glucose concentration signals (metabolic hypothesis). However, recent experimental evidence supports a number of pathways that involve release of vasoactive messengers (neurogenic hypothesis), which are displayed in Fig.1.2.

A number of labs now propose that both of these mechanisms occur [187][161], explaining how CBF can increase within 1s (supporting the metabolic hypothesis as messenger molecules do not reach arteries in that time [85][171]), why CBF increases overcompensate for energy usage (supporting the neurogenic hypothesis, which does not have negative feedback[c] [8]), and why negative BOLD signals have been observed in ipsilateral cortex (also supporting the neurogenic hypothesis as this drop in CBF is not matched by energy changes [44]).

[a] Cellular respiration creates the energy (ATP) that cells use to do work (such as carry and transmit electrical signals in the case of neurons) from glucose and O_2. Glycolysis, which is the first stage of cellular respiration, is detailed in Section 3.1.2.1.

[b] Some recent evidence suggests that, despite there being no smooth muscle in their walls, capillaries may also moderate CBF, mediated by pericytes [149]

[c] In the metabolic hypothesis, because energy demand causes CBF increases, only enough CBF to fuel the energy demand is delivered.

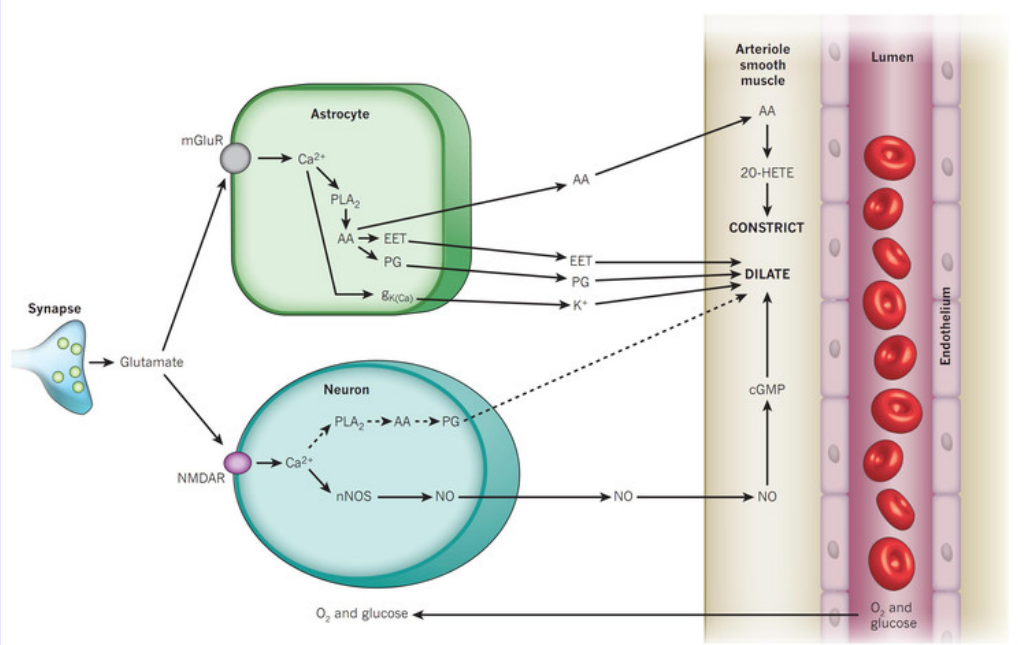

Figure 1.2: Postulated neurogenic neurovascular coupling signalling pathways. Glutamate released from pre-synaptic neurons triggers calcium waves in both neurons and astrocytes. In neurons, the heightened calcium levels lead to production of Nitric Oxide (NO) and Prostaglandin (PG), causing vasodilation in arterial smooth muscles. In astrocytes, the heightened calcium levels lead to production of arachidonic acid, which causes both vasodilation via epoxyeicosatrienoic acids (EETs) and PG, and vasoconstriction. Potassium (K^+) dynamics in astrocytic endfeet are also vasodilatory. Figure reproduced from Atwell et al., 2010, Figure 2 [8].

(electrophysiology) and hemodynamics (BOLD signal or a proxy measurement) and then compare their linear relationships across groups. These neurovascular coupling studies fall into two groups:

1. Studies aiming to validate BOLD GLM analyses by investigating the parametric relationships between neuronal activity and hemodynamics (the parameters typically investigated are neuronal and hemodynamic response amplitude and/or power). These studies usually fall into two sub-groups: 1) those that modulate some aspect of the stimulus paradigm (example studies include Ngai et al., 1999, [126] and Ureshi et al., 2003 [190]), and 2) those that compare control and neuromodulated groups (an example study is Franceschini et al., 2010 [56]).

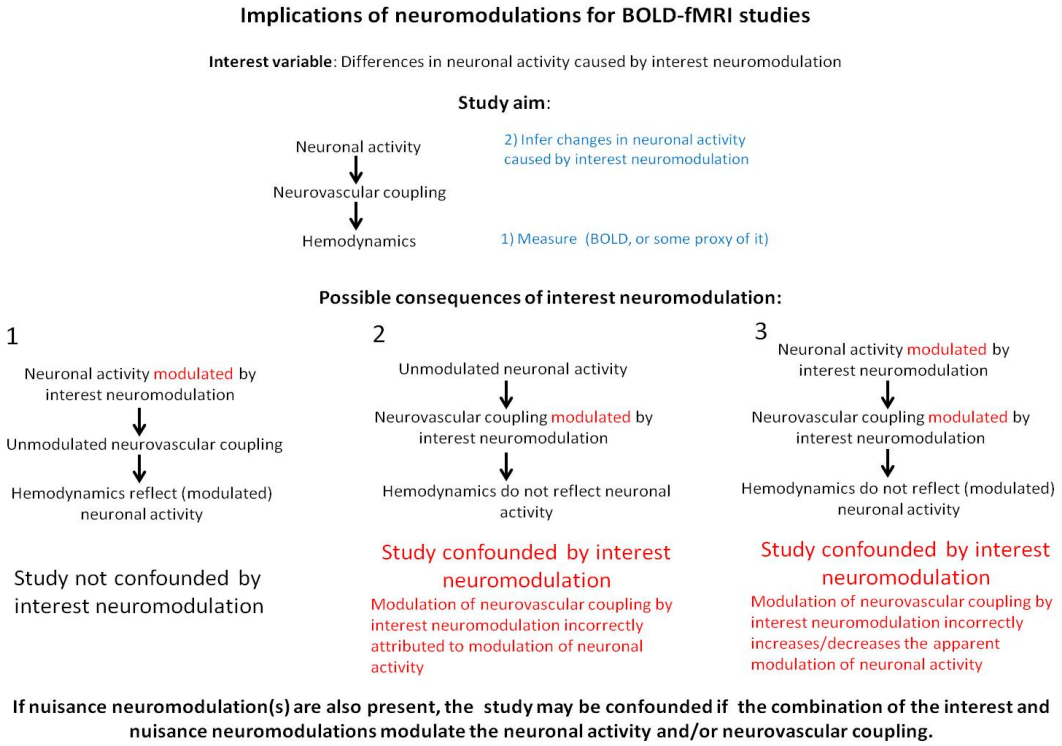

Figure 1.3: It is important that interest neuromodulations do not modulate aspects of neurovascular coupling of relevance to BOLD signal interpretation because otherwise these effects will sum with the effects of the interest neuromodulation on neuronal activity, making it appear that the interest neuromodulation alters neuronal activity more or less than it actually does. Moreover, it is important to appreciate that many BOLD experiments will have more than one neuromodulatory factor. This is because nuisance neuromodulations, i.e. neuromodulations that do not separate the experimental groups, are unavoidable for some BOLD studies: nuisance neuromodulations relevant to human BOLD studies include aging [42] and substances such as caffeine [121], however anaesthesia is a pre-requisite nuisance neuromodulation for most preclinical BOLD studies. When nuisance neuromodulations are unavoidable, it is the confounding effect caused by the combination of interest and nuisance neuromodulations that needs to be considered.

2. Studies concerned with identifying the processes involved in neurovascular coupling by pharmacologically modulating a theorised component (an example study is Mathiesen et al., 1998 [114]).

For parametric neurovascular coupling studies, the null hypothesis usually being tested by the traditional approach (linear regression of multimodal shape features) is that the linear relationship between the neuronal activity and hemodynamics is the same for different stimulation paradigms/neuromodulations, and therefore that BOLD GLMs are not confounded by using different stimulation paradigms/using the same GLM for both the control and neuromodulated groups. Similarly, for neurovascular coupling studies aiming to elucidate the role of some biological process in neurovascular coupling, the null hypothesis usually being tested is that the regression coefficients are the same for the control and neuromodulated groups, and therefore that the interest process does not impact on neurovascular coupling.

This traditional analysis has gained popularity because it is generally accepted that there is an approximate linear relationship between summed neuronal activity and the magnitude of the accompanying hemodynamic response [16][104]. However, the traditional analysis does not allow the researcher to make any robust conclusions about the implications of neuromodulations for GLMs as it is the standard error of the model fit as well as its regression coefficient magnitude that determines whether the BOLD response is detected as significant or not, and the traditional analysis does not consider the standard error of the model fit. What is more, shape features do not have clear biological interpretations, meaning that knowledge of a disease or substance's actions on neural biology cannot be exploited when developing strategies to account for neurovascular coupling changes in the GLM.

Because of these limitations, there is a growing consensus that studies aiming to investigate aspects of neurovascular coupling should move away from being correlation driven, such as linear regression of shape features, to being mechanistic,

i.e. drawing on our understanding of the biological processes facilitating neurovascular coupling [45]. The overarching goal of this thesis, therefore, is to:

"describe, demonstrate and evaluate the novel use of a mechanistic method based on generative embedding for making comprehensive, relevant and interpretable conclusions about the implications of neurovascular coupling modulations for BOLD imaging".

Brodersen et al., 2011, were the first to coin the term generative embedding for describing a framework for statistically interrogating the hidden physical processes generating observable neuroimaging data [19]. In their work, Brodersen et al. classified subject groups by embedding the posterior means from fitting generative models to (separately acquired) BOLD-fMRI and local field potential (LFP - synaptic (low-frequency) neuronal activity between neurons recorded using electrodes) signals in support vector machines. They argued that classifying a small set of biologically meaningful parameters rather than the neuroimaging timeseries themselves or their shape features not only improved the performance of their classifiers because the informative features of the data were no longer obscured by the uninformative features, but also provided results with clearer mechanistic interpretations [18][19].

1.2 Thesis Summary: Using generative embedding to assess the effects of neuromodulations on BOLD-fMRI signals

All of the generative embedding algorithms used in this thesis consist of the same three steps:

1. **Generative** modelling of neuroimaging signals
2. **Embedding** of the generative model posterior parameters into statistics

3. Inference on a hypothesis using these statistics

In this thesis, generative embedding is used to test hypotheses about the relationships between the biological processes 1) underlying the generation of evoked neuronal responses and 2) effecting neurovascular coupling. Box 3 attempts to justify this approach to investigating the consequences of neuromodulations for neurovascular coupling.

The inferences made in this thesis fall into two main groups:

1. Chapters 2 to 4: Inferences that test whether neuromodulations of neurovascular coupling confound BOLD GLM analyses,
2. Chapters 5 to 7: Inferences that explore whether neuromodulations of neurovascular coupling can act as tools for clinical practice and/or drug development.

1.2.1 Neuromodulations as confounding factors for BOLD-fMRI studies

1.2.1.1 Chapter 2

To demonstrate the utility of the generative embedding framework for quantifying the extent to which a neuromodulation confounds BOLD GLMs (outlined in Fig.1.4), Chapter 2 describes the GLM analysis of the preclinical anaesthetics dataset collected by Franceschini et al.[1] [56] (Step 1.1 of the generative embedding framework). The traditional intepretation of this GLM analysis is that anaesthetics modulate neuronal activity to different extents, resulting in different neural activation maps for the different anaesthetics. Given the evidence in the literature that anaesthetics act differently on vascular as well as neuronal processes, I hypothesise that differences in neurovascular coupling caused by the interest neuromodulations are also responsible for the different GLM significances.

[1]This dataset has only one neuromodulatory factor: the interest neuromodulation of anaesthetic agent.

Box 3: Evidence for the coupling of processes underlying evoked neuronal responses and effecting neurovascular coupling

The major assumption behind using GLM to predict neuronal activity from hemodynamics is that hemodynamics are proportional to the local average neuronal activity. The generative embedding approach taken throughout this thesis to investigate the consequences of neuromodulations for neurovascular coupling makes a deeper assumption, namely that the biological processes underlying the generation of evoked neuronal responses are coupled to the biological processes that effect neurovascular coupling.

There is considerable evidence in the literature that supports this assumption. In the case of the metabolic hypothesis, inhibitory neurotransmitter adenosine has been observed to increase low-frequency EEG power and to shift certain narrow frequency bands to lower values and/or increase their power [64]. It has also been observed to function as a metabolic messenger of neural energy demand that gives rise to changes in CBF [88]. Similarly, in the case of the neurogenic hypothesis, experimental evidence suggests that the excitatory neurotransmitter glutamate released during neuronal processing can not only be reliably detected from evoked neural response generative model parameters [118], but it also causes the release of calcium that leads to vasodilation [23]. By being causally related to both evoked neuronal responses and hemodynamic changes, adenosine signalling and glutamate release mediate the covariation of the processes underlying evoked neuronal responses and the processes effecting neurovascular coupling.

It is worth mentioning that the generative embedding approach taken throughout this thesis to investigate the consequences of neuromodulations for neurovascular coupling is made possible because all biological processes are stochastic, meaning that neuroimaging signals exhibit small trial-by-trial intra-subject variations from which the relationships of their underlying biological processes can be estimated.

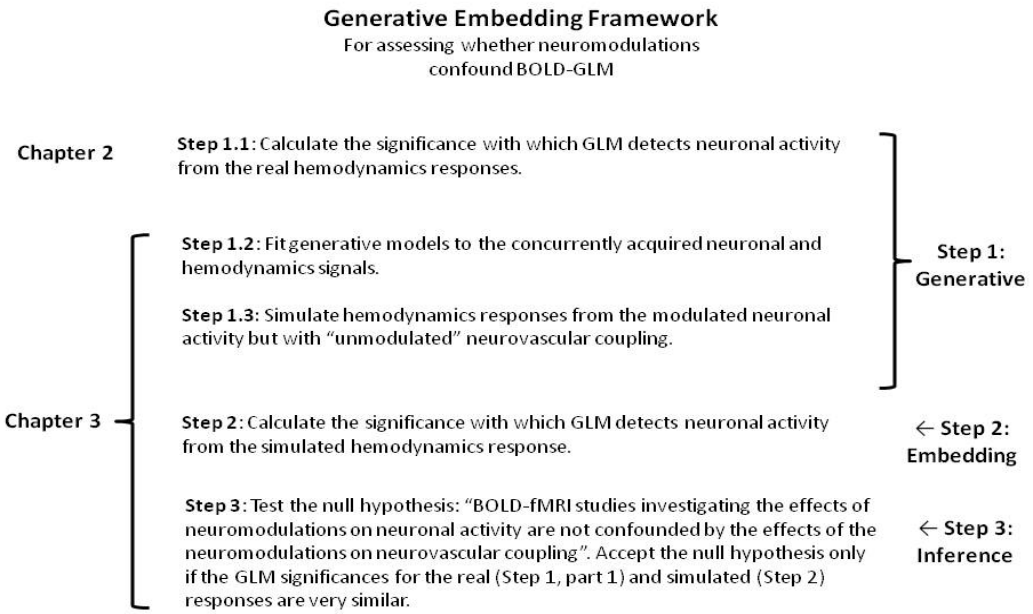

Figure 1.4: Overall scheme of the novel generative embedding framework presented in Chapters 2 and 3 for understanding whether neuromodulations confound BOLD-fMRI studies analysed using standard GLM.

1.2.1.2 Chapter 3

The aim of Chapter 3 is to test the hypothesis that it is both the differences in neurovascular coupling as well as the differences in neuronal activity that are responsible for the different GLM significances across anaesthetics. Step 1.2 of the generative embedding algorithm involves the fitting of a generative model to the concurrently acquired neuronal and hemodynamic signals in order to summarise the signals using a small number of biologically relevant parameters. These posterior parameter estimates are then used to simulate hemodynamic signals in which differences in neuronal activity exist only, and GLMs are fitted to these simulated signals (Step 1.3). Step 2 of this generative embedding procedure is the embedding of these GLMs into analysis of variance statistics. By comparing the GLM significances of the real and simulated data, it is possible to quantify the extent to which the effects of the anaesthetics on neurovascular coupling caused the

differences in GLM significances across the anaesthetics in the real data (Step 3).

1.2.1.3 Chapter 4

There are two interdependent ways in which neuromodulations can confound the commonly used BOLD GLM:

1. By altering the linearity assumption of the hemodynamic response that is assumed to exist in unmodulated brain,
2. By altering the hemodynamic response to a single, short stimulus, so that its basis functions, which were optimised for unmodulated brain, are no longer appropriate.

Chapter 4 begins by showing the extents to which these two phenomena cause GLM analyses to be confounded under each anaesthetic. To gain some insight into which biological processes within the neurovascular coupling mechanism may be responsible for these confounding phenomena, a different generative embedding algorithm is used: the posterior estimates from fitting the generative model to the real data are embedded into regression statistics in order to identify which ones significantly explain the degrees to which the different anaesthetics confound the GLM analyses. The biological insight gained from this chapter has the potential to make existing algorithms that reduce the inaccuracies of BOLD GLM analyses caused by neuromodulations more effective and/or practical to use without complicating the standard (ANOVA) way of reporting BOLD results.

1.2.2 Neuromodulations as tools for neuroimaging investigations

1.2.2.1 Chapter 5

A technological solution for acquiring concurrent measures of neuronal activity and hemodynamics and then detecting brain pathologies from the inferred

neurovascular coupling has been developed by Dartmouth Brain Imaging Group [67]. The algorithm used in the software of this prototype product is based on the traditional analysis, i.e. linear classification of the multimodal shape parameter space. To provide some understanding of the ability of this system to classify healthy and diseased brains, Chapter 5 describes the collection of a preliminary preclinical dataset consisting of the concurrently imaged neuronal activities and hemodynamics of healthy animals and animals in which a neurodegenerative disease-like state has been experimentally induced. By analysing these data using the traditional analysis approach, it is shown that there are some significant differences in the linear relationships of the healthy and diseased neuroimaging signals for the prototype solution to use for disease classifications.

1.2.2.2 Chapter 6

As already demonstrated in the earlier chapters, generative embedding algorithms provide a much more comprehensive and informative understanding of neurovascular coupling than the traditional analysis. The main aim of Chapter 6 is to investigate the capabilities of generative embedding algorithms for commercial products, such as the prototype product already mentioned, that make some inference about brain disease state by probing aspects of neurovascular coupling. Depending on which statistics the posterior parameters are embedded into, these inferences span diagnosis of brain disease state (in line with the prototype product), to monitoring of specific aspects of neurobiology during disease progression, to detailed understanding of the biological actions of treatments such as prescription drugs. However, what all of these applications of generative embedding algorithms have in common is that they benefit from increased discriminability of the healthy and diseased posterior parameter spaces. To quantify the upper limits of the discriminability of the healthy and neurodegenerative-diseased posterior parameter spaces, the posterior parameters

from the fitting of generative models to the dataset collected by Rosengarten et al., 2006, are embedded into mutual information statistics. Because mutual information is sensitive to all of the information in the posterior parameter space, this provides an upper limit on the discriminability of the healthy/diseased parameter spaces that might be used to make some inference about brain disease state or treatment effect.

The second aim of Chapter 6 is to compare the discriminabilities of the generative model parameter spaces when the two imaging signals are acquired concurrently (exploiting the differences in neurovascular couplings between the diseased and healthy animals), to when they are acquired separately (when the differences in neurovascular couplings are not captured). By showing that the healthy/diseased posterior parameter spaces are more discriminable when the neuronal activity and hemodynamics have been imaged concurrently, the case for developing hardware similar to the prototype product that is capable of measuring neuronal activity and hemodynamics concurrently and is supported by generative embedding algorithms is strengthened.

1.2.2.3 Chapter 7

Chapter 7 considers the commercial potential of these novel technologies for investigating neurovascular coupling using generative embedding algorithms. The two healthcare markets that could benefit from this software are:

1. Clinical, for improved diagnosis and monitoring of neurodegenerative diseases,
2. Pharmaceutical, as a tool for the development of drugs for neurodegenerative diseases.

1.3 Novel contributions to the field

This thesis describes the first uses of generative embedding for investigating the implications of neurovascular coupling modulations on BOLD imaging. The most important novel contributions of this thesis enabled by these generative embedding algorithms are:

- **Chapters 2 and 3: Precise quantification of the extent to which neuromodulations confound BOLD GLM analyses**

  Until now, experimental methods such as comparison of hemodynamic responses to CO_2 challenge (to assess cerebrovascular reactivity) [202] have been used to identify whether neuromodulation confound BOLD GLM analyses. No reference can be found in the literature about the implications of these confounding effects on neuronal activation statistical significances being explicitly quantified. Therefore, these chapters describe the first known attempt to precisely quantify the extent to which neuromodulations confound BOLD GLM analyses and it boasts the added advantage of not requiring any additional experimental data to be collected.

- **Chapter 3: Variational Bayes fitting of a biophysical model to concurrent measures of neuronal activity and hemodynamics in which the multimodal parameter correlations are estimated**

  A small number of journal papers have published attempts to fit biophysical models of hemodynamics driven by neuronal activity to concurrent measures of neuronal activity and hemodynamics (notably Rosa et al. [163]). However, because they first optimise the model fit to the neuronal activity and then optimise the model fit to the hemodynamics afterwards, the covariances of the multimodal parameters cannot be estimated. In this chapter, estimation of these multimodal parameter covariances is achieved by iterating between 1) fitting the neuronal model to

the neuronal activity, and 2) fitting the full model (neuronal activity → metabolism → hemodynamics) to the hemodynamics data.

- **Chapter 4: Biologically interpretable insight into how best to compensate for the confounding effects in BOLD GLM analyses**

  Although attempts to explain the extent to which neuromodulations confound BOLD GLM analyses using shape features have already been published [72][101], this chapter describes novel work because it enables biologically interpretable conclusions to be made about the actions of neuromodulations that confound BOLD GLM analyses, informing how best to modify the standard GLM to compensate for these confounding effects and increasing the effectiveness and practicality of established correction algorithms.

- **Chapter 6: Comprehensive quantitative assessment of the value of hybrid neuroimaging technologies for diagnosing, monitoring and understanding neurodegenerative diseases**

  As already mentioned, a prototype hybrid neuroimaging solution that use the traditional analysis method (linear regression of the multimodal shape features) to make inferences about neurodegenerative disease state has already been developed. This chapter describes the first attempts to understand the limits of this solution for providing information about neurodegenerative disease state if generative embedding were used instead of linear regression of shape features. This has involved the derivation of a new, less-biased method for calculating the mutual information and synergy of spaces involving categorical variables using k-nearest neighbours.

Chapter 2

Testing BOLD-fMRI general linear model sensitivity to anaesthetic neuromodulation

The primary aim of Chapters 2 to 4 of this thesis is to demonstrate the utility of generative embedding for investigating the extent to which neuromodulations confound BOLD-fMRI experiments analysed using standard GLM. This aim is motivated by the fact that the traditional approach for investigating the effects of neuromodulations on BOLD GLM analyses can only offer limited insight about the extent to which neuromodulations confound BOLD studies and through which biological actions. This is because the traditional approach only considers the effects of neuromodulations on the linear relationships between the shape features of neuronal activity and hemodynamic outputs. In contrast, by being sensitive to all orders of correlation between a comprehensive set of mechanistically interpretable parameters, generative embedding can provide information about the confounding effects of neuromodulations that can be used to compensate for them, improving the accuracy of any conclusions made from BOLD-fMRI studies.

This chapter describes the implementation of the first step of a generative

embedding algorithm used to quantify the extent to which neuromodulations confound BOLD GLMs. Specifically, the significances with which the standard BOLD GLM detects neuronal activity from the hemodynamics responses recorded under a range of neuromodulations are calculated.

As it would be unfeasible to collect a dataset of sufficient size and range of neuromodulatory effects during this DPhil, and because one aim of scientific research is to reduce the number of animals used in experiments, a suitable, previously published dataset was requested and received from Franceschini et al., 2010, for this work [56]. This dataset comprises measures of neuronal activity and hemodynamics concurrently recorded under six different anaesthetic modulations.

2.1 Anaesthetics dataset

The aim of the original presentation of Franceschini et al.'s dataset was to understand whether anaesthetics confound parametric neurovascular coupling studies (whose aim it is to validate BOLD GLM analyses using the traditional method). This work is important because neurovascular coupling experiments are usually carried out on laboratory animals, as that enables invasive imaging and therefore collection of a less noisy and more informative dataset. Unless animals are trained to lie still [110], which is a time-consuming (and therefore costly) exercise that very few research groups opt for, they need to be anaesthetised to minimise head motion and animal stress. Unfortunately, different anaesthetic protocols (anaesthetic agent(s), dose and administration route) are used across research groups, meaning that incorrect conclusions may be drawn about neurovascular coupling from the comparison of the conclusions from different experiments if anaesthetics modulate neurovascular coupling in different ways.

Franceschini et al. tested the null hypothesis that anaesthetics do not confound parametric neurovascular coupling studies using the traditional analysis approach (linear regression of the multimodal shape features). They found that, using the

same linear regression equation for all six anaesthetics, they were able to predict all of the hemodynamic shape features (areas under the main hemodynamic response) from the corresponding neuronal shape features (areas under the SEP P1, N1, P2 and N2 features - see Fig.2.1) with a correlation coefficient of 0.96 (see Equ.2.1). This maintainance of the linear relationships between the multimodal shape features across all six anaesthetics indicates that anaesthetics do not confound parametric neurovascular coupling studies.

$$Hb = \sum_{i=1}^{4} (h_i \cdot SEP_i), \tag{2.1}$$

where Hb = hemodynamic shape feature, SEP_i is the i^{th} of the four neuronal activity shape features and h_i is its regression coefficient.

The traditional analysis was sufficient for Franceschini et al.'s work because the studies to which that work is pertinent, namely parametric neurovascular coupling studies, are insensitive to: 1) the nonlinear aspects of neurovascular coupling, and 2) modulations of the processes that control the complete shape (rather than just the power) of hemodynamic responses. However, when instead considering the implications of neuromodulations for BOLD-fMRI studies analysed using GLM, the traditional analysis is no longer sufficient because GLM statistics are a function of GLM fit, which is determined by the shape of the hemodynamics responses (and not inferred in parametric neurovascular coupling studies), as well as the amplitude/power of the hemodynamics (the shape features typically extracted for traditional analyses). To gain a comprehensive understanding of whether anaesthetics confound BOLD-fMRI studies analysed using GLM, it is necessary to consider their impacts on both the linear and nonlinear aspects of neurovascular coupling that determine the hemodynamic response shape and

amplitude/power, and this can be achieved using generative embedding[1].

2.1.1 Experimental protocol

The neuronal and hemodynamic responses of 33 male Sprague-Dawley rats to parametric forepaw electrical stimulation were concurrently recorded using the following imaging modalities, respectively:

- Scalp electroencephalography (EEG), which measures the aggregate neural electrical potential from the activity of millions of neurons below the recording site. This was used to record the somatosensory evoked potentials (SEPs, the neural electrical responses to sensory stimuli).
- Diffuse optical imaging (DOI), which measures absorbance of laser light at two different wavelengths to calculate changes in oxyhemoglobin/deoxyhemoglobin (HbO/HbR) concentration. This imaging modality was used to record changes in concentrations of HbR ([ΔHbR], a proxy for the BOLD signal). One advantage of working with DOI data rather than the BOLD signal itself is that DOI boasts much greater temporal resolution, which will prove useful when fitting generative models to these timeseries in the next two chapters.

All rats first received Isoflurane (2 - 2.5%) administered via a face mask in a gas mixture of 80% air and 20% oxygen to induce anaesthesia. They then received one of the following anaesthetics (neuromodulations) during the experiment: Alpha-Chloralose (40mg/kg/h intravenous (i.v.), six rats), Pentobarbital (25mg/kg intra-peritoneal (i.p.), five rats), Ketamine (20mg/kg/h - 2mg/kg/h i.v., five rats, delivered in combination with sedative Xylazine), Fentanyl (90g/kg/h - 4.5mg/kg/h i.v, six rats, delivered in combination with neuroleptic Droperidol),

[1]It is appreciated that anaesthesia is an unusual interest neuromodulation for a BOLD study, however the main aim of its use in Chapters 2 to 4 is only to demonstrate the ability of the novel generative embedding framework for understanding whether neuromodulations confound BOLD-fMRI studies analysed using GLM.

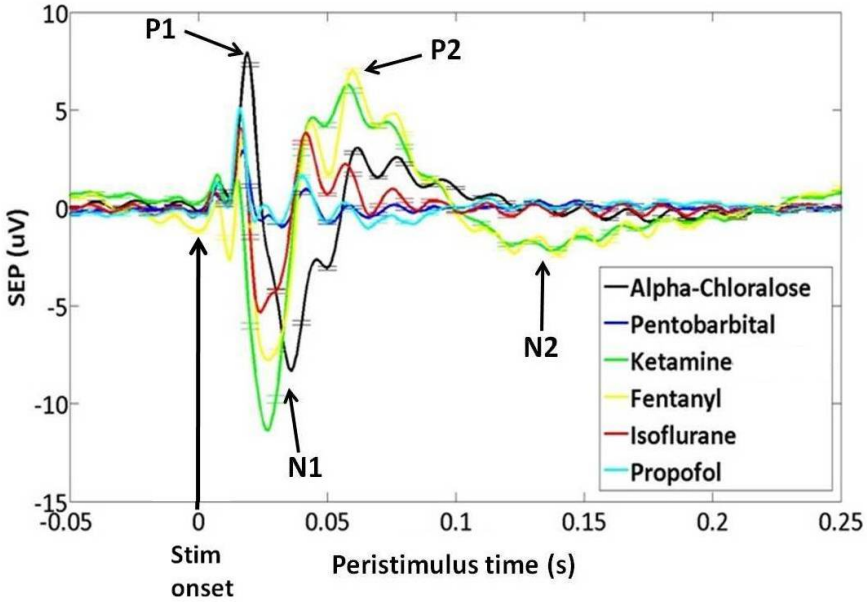

Figure 2.1: Mean SEP responses to a single stimulus under the six anaesthetics. Errors are standard errors. P1, N1, P2 and N2 shape features are highlighted.

Isoflurane (1.2% gas, five rats) or Propofol (50mg/kg/h i.v., six rats). These anaesthetic agents, doses and administration routes are commonly used in preclinical BOLD studies, and the anaesthetics cover a wide range of anaesthetic types, including volatile (Isoflurane), barbituate (Pentobarbital) and opioid (Fentanyl).

On average, each rat underwent five runs of stimulation in each forepaw. These runs comprised trains of 200s stimuli delivered at 3Hz with pseudo-randomly arranged lengths of 1s, 3s, 5s, 7s, 9s, 11s or 13s and an average inter-stimulus interval (time between the starts of subsequent trains) of 12s. Each train duration was delivered on average five times in each run.

The mean EEG and [ΔHbR] responses to all trials of 5s of stimulation across all animals under each anaesthetic are shown in Fig.2.1 and Fig.2.2, respectively. The EEG responses comprise four features: the initial, fast peak (P1), then a slower trough (N2), then a second, still slower peak (P2), and, for some anaesthetics, a second, yet again slower trough (N2). The origins of these features are described in Fig.2.3.

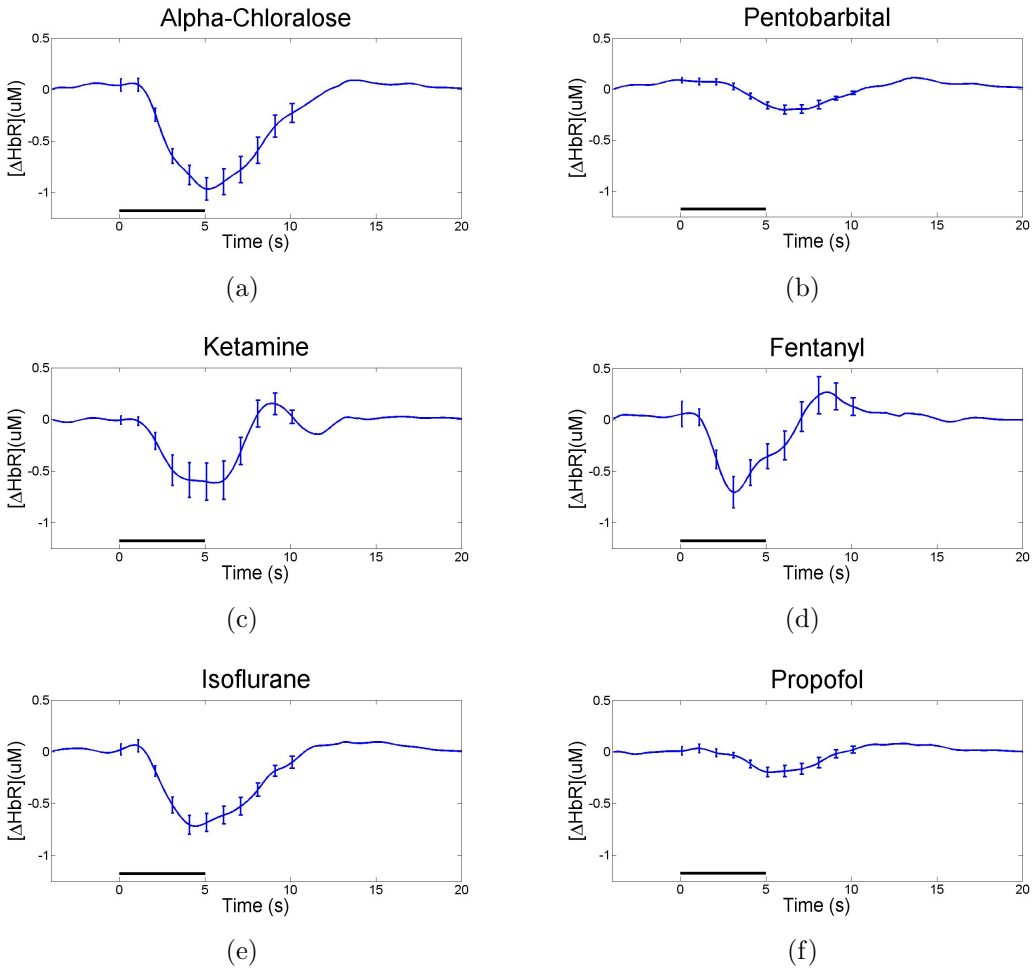

Figure 2.2: Mean [ΔHbR] response to 5s trains of the stimuli under the six anaesthetics. As stimuli were often delivered before the response to a previous stimulus had returned to baseline (the average inter-stimulus interval is only 12s), only the response occuring before the next stimulus was delivered were used to calculate these mean responses. Errors are standard errors and are only included for the post-stimulus period equal to the shortest individual inter-stimulus interval that occurred under that anaesthetic.

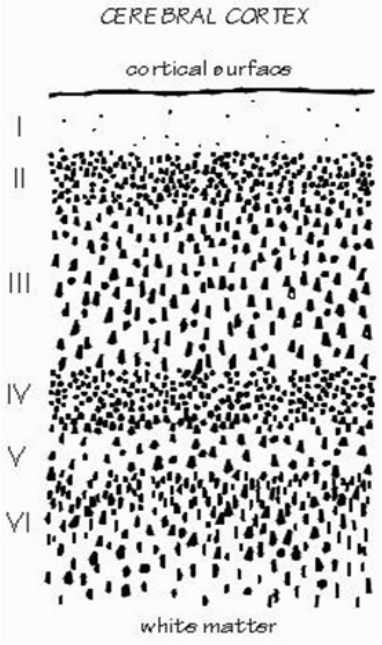

Figure 2.3: Anatomy of a cortical column. Thalamic input arrives in layer IV, the granular layer, causing a current sink there and therefore a current source in the superficial layers. Because it is the activity of the superficial layers that is most strongly detected by EEG, the current source causes P1 in the EEG signal. The pyramidal cells in the supragranular layers (layers I and II) are now activated, causing a current sink to form there and giving rise to N1 in the EEG signal. Finally, cortico-cortical projections from these superficial layers transmit electrical responses to other cortical columns in brain regions that do not receive electrical impulses directly, giving rise to P2 and N2 in the SEP [157]. Figure reproduced from Calvin and Ojemann, 1994, Figure 11 [28]

2.2 GLM analysis

GLM was established as the standard method for statistically comparing fMRI signals in 1995 [60]. Since then, improvements and extensions have been incorporated into its implementations in the major fMRI analysis software packages but its basic utility remains unchanged[1]: create an expected BOLD response to the known pattern of stimulation using linear modelling and then compare it to the observed BOLD responses. Significant BOLD responses are those that match the expected BOLD response with highly positive (for neural activation) or negative (for neural deactivation) correlation, i.e. with large size and goodness of fit. The assumption is that the locations of these significant BOLD responses correspond to the locations of significant neural processing, i.e. the neural processing of interest to the BOLD study has generated the detected significant BOLD response not only in the same location in the brain, but with a proportional magnitude.

[1] Although arguably more accurate methods have since been developed for fMRI analyses, the conceptual simplicity of GLM has secured its longevity as the standard analysis tool: it only involves linear regression and analysis of variance (ANOVA), both of which are usually covered in undergraduate science degrees.

The expected response is created from a linear function of explanatory variables, each of which aims to model one aspect of the observed BOLD signal, such as the response to one stimulation type within the stimulation paradigm, or Gaussian noise. The explanatory variables that aim to model the responses to one type of stimulation are formed by convolving delta functions mirroring stimulation timings with a number of basis functions, which may include 1) a canonical hemodynamic response function (HRF), 2) the first temporal derivative of this canonical HRF, and 3) the dispersion derivative of this canonical HRF. The purpose of these derivative basis functions is to allow the GLM to be more flexible with regard to neurovascular coupling. For example, the temporal derivative allows the canonical HRF to shift in time so that the onset of an expected BOLD response can better match the onset of the observed BOLD response it is fitting to. The dispersion derivative allows the width of the canonical HRF to vary, providing a way for the shape of the expected BOLD response to better match the shape of the observed BOLD response too. Therefore, if any neuromodulation alters neurovascular coupling in a way that causes changes to the onset and/or temporal dynamics of the hemodynamic responses, inclusion of the respective derivatives as basis functions in the design matrix of explanatory variables may improve GLM fits considerably, helping to unmask the true significance with which a brain region is involved in processing the stimulation paradigm.

The general claim of BOLD GLM is that it infers the amplitude of neuronal activity by drawing on the assumption that the amplitude (firing rate) of the neuronal activity is proportional to the amplitude (magnitude) of the hemodynamic response. However, it is possible (and likely) that relationships exist between the various aspects of neuronal activity and hemodynamics responses, as shown in Fig.2.4. Therefore, although the standard interpretation for a BOLD GLM analysis that detects neuronal activity with different significances under different neuromodulations would be that the amplitudes of the neuronal activity

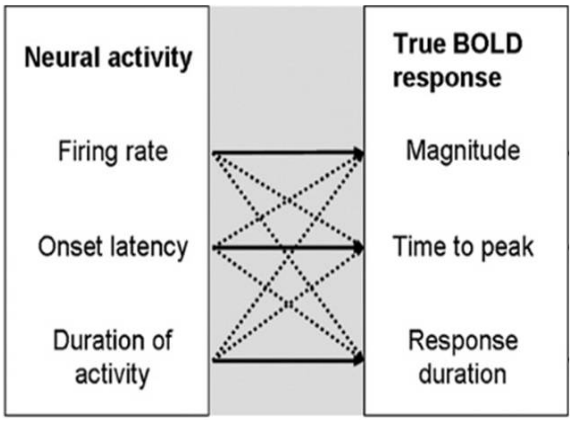

Figure 2.4: Hypothesised relationships between three selected shape features extracted from measures of neural activity and three selected shape features extracted from evoked changes in the BOLD response. Solid lines indicate relationships that are likely to be relatively direct, including the relationship between neuronal activity firing rate and BOLD magnitude, which forms the basis of GLM analyses for detecting neuronal activations from BOLD signals. Dashed lines indicate relationships that are likely to be more indirect, including relationships between neuronal activity onset latency and hemodynamic mangnitude, and between neuronal activity duration of activity and hemodynamic magnitude, neither of which is formally considered by GLM. Figure reproduced from Lindquist, 2009, Figure 1 [100]

are modulated to different extents, in reality the differences may actually be attributable to other aspects of neuronal activity, such as neuronal activity onset latency and duration of activity, and/or to other differences in the hemodynamic response that impact on the fit of the observed hemodynamics to the expected hemodynamics, such as hemodynamic time to peak and response duration.

The complexity of BOLD GLM interpretation highlights the need for a rigorous approach, such as generative embedding, to investigate the consequences that neuromodulations have for this widespread method of analysis.

2.3 Methods

In this work, the MATLAB (2012b, The MathWorks, Inc., Natick, Massachusetts, United States) GLM implementation in SPM8[1] (http://www.fil.ion.ucl.ac.uk/spm/) was adapted to run the standard

[1]Statistical Parametric Mapping (SPM) is the most widely used software for fMRI analysis and visualisation.

multi-subject GLM analysis of the [ΔHbR] data, resulting in two neural activation significance levels per rat (one for each hemisphere of brain recorded from). The details of this standard GLM implementation are presented in Appendix A[1]. Specifically, four sets of GLMs were fitted to the [ΔHbR] signals: once using only the awake-rat [ΔHbR] equivalent of the standard canonical HRF as a basis function for the regressor in the GLM, and then again with the temporal and/or dispersion derivatives as additional basis functions.

The canonical HRF used by default in SPM8 is a double gamma function that Friston et al., 1998, fitted to the average BOLD response recorded during a single awake-human fMRI session [62]. As the aim of this section of thesis is to demonstrate the potential of generative embedding for quantifying the extent to which a neuromodulation confounds BOLD GLM analysis using the case study of anaesthetic neuromodulation of rat BOLD signals, it was necessary to generate the awake-rat [ΔHbR] equivalent of canonical HRF (from now on referred to simply as canonical [ΔHbR]). This was achieved by deconvolving the average HbR response to 2s of 1Hz whisker stimulation in awake rat recorded in Martin et al.'s 2006 study [110], resulting in the canonical [ΔHbR] shown in Fig.2.5 after smoothing.

Finally, in order to make general statements about the extent to which the different anaesthetics confound GLM analyses, second-level random effects analysis was used to calculate the anaesthetic-specific multi-subject ("group-level") significances. The reliabilities of the GLMs to detect significant neural activation under the different anaesthetics were then found by calculating the percentages of rat brain hemispheres in which significant neural activations were detected.

[1]Although the GLM implementations in SPM and other neuroimaging software packages, such as FSL, ANFI and BrainVoyager, have been found to give similar GLM statistics [119], there are differences in their implementations. Some of these differences are described in Appendix A.

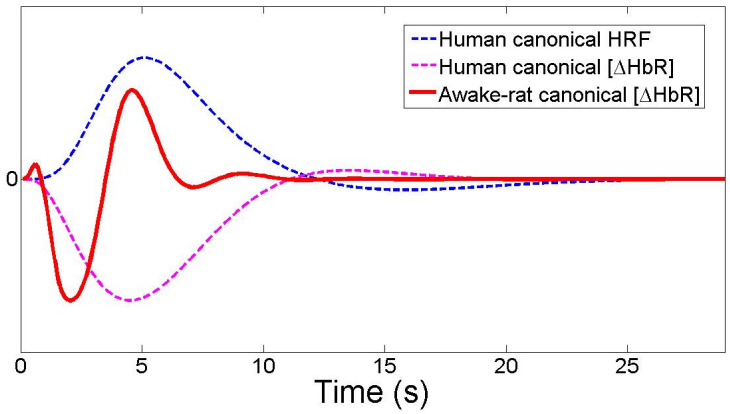

Figure 2.5: Awake-rat canonical [ΔHbR] used as a basis function for fitting GLMs to [ΔHbR] signals. The human canonical HRF and human canonical [ΔHbR] (simulated by fitting DCM-fMRI, which is described in detail in Chapter 3, to canonical HRF) are also plotted to enable comparison (scaled to enable clearer comparison of shape).

2.4 Results and Conclusions

Fig.2.6 compares the mean [ΔHbR] responses under each of the six anaesthetics with the canonical [ΔHbR], which is scaled so that they have equal amplitudes. Qualitative descriptions of the ways in which these responses differ from the canonical [ΔHbR] are listed in Table 2.1, along with the expected consequences of these differences for the GLM analyses. These expected consequences are based on the **prediction** that the anaesthetic modulatory effects on hemodynamics typically result in Z-values being underestimated, with this underestimation being reduced as the [ΔHbR] derivatives are added as basis functions because they reduce the model fit standard error[1].

The results from the GLM analyses are shown in Tables 2.2 and 2.3. Specifically, the significances and reliabilities with which the GLM framework employed by SPM8 can detect neural activations in anaesthetised rats are reported. Sample GLM fits for each of the anaesthetics are shown in Fig.2.7 - 2.12

[1]Group-level GLM Z-values, such as those calculated for each anaesthetic later in this chapter, are a function of the mean and variance of the Z-values calculated for each individual subject. Because the predictions made in Table 2.1 are based on the mean responses under each anaesthetic only, they do not take the variance in the size and shape of the responses into account. Therefore, if these predictions turn out to be grossly incorrect, it will be because of this variance.

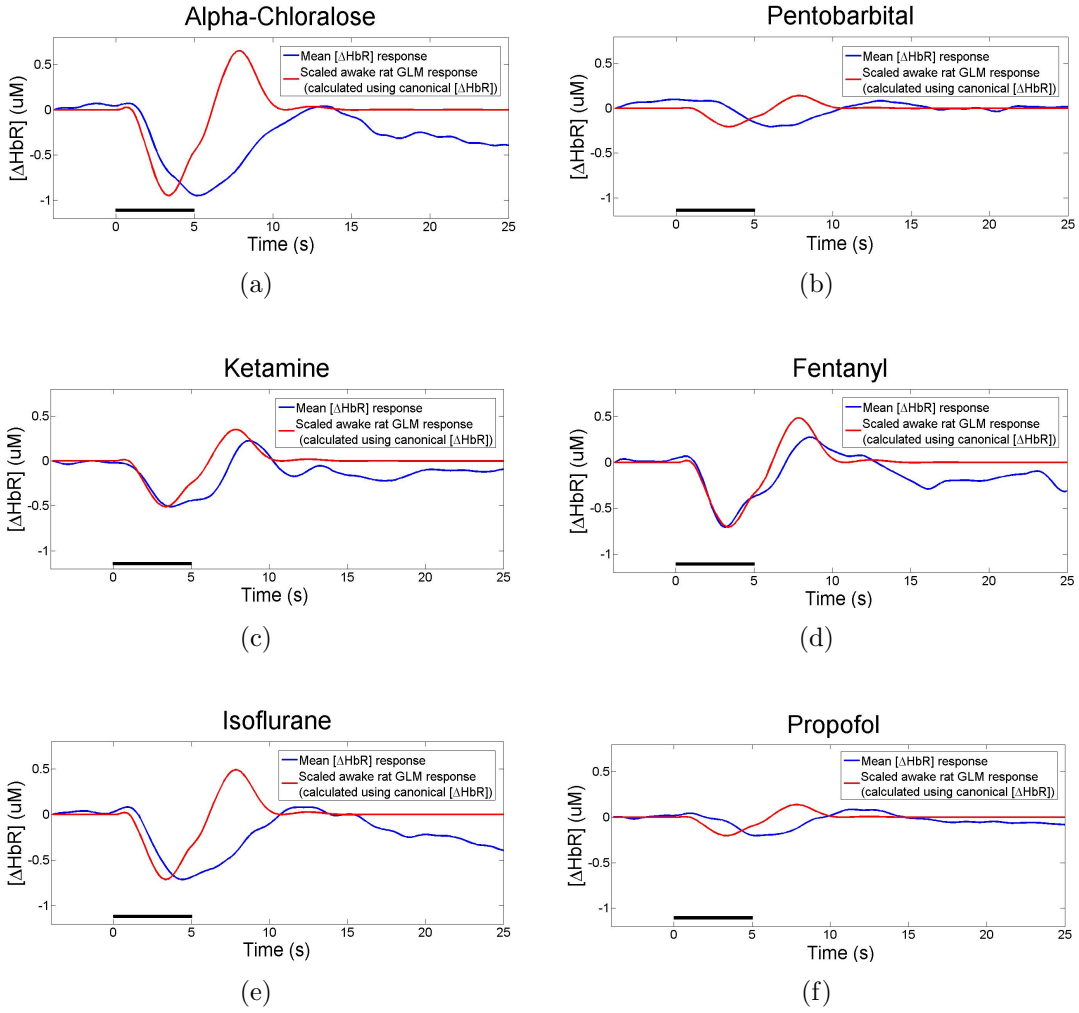

Figure 2.6: Comparison of the mean [ΔHbR] responses to 5s trains of stimuli under each of the six anaesthetics (black) and the response simulated by convolving canonical [ΔHbR] with this same stimulus train (red).

(these samples are taken from the rats found to have the highest individual Z-values under each anaesthetic).

These statistical results are in line with the qualitative observations listed in Table 2.1 in the following ways:

- Fentanyl has a high Z-value, Alpha-Chloralose has a medium Z-value and Pentobarbital has a low Z-value when only the canonical [ΔHbR] is included.
- Adding the dispersion derivative improves Z-values under Alpha-Chloralose, Isoflurane and Propofol, and to a small degree under Pentobarbital, Ketamine and Fentanyl too.

Anaesthetic	**Response size**	**Response shape**	**Response onset**	**Response width**	**Canonical [ΔHbR] Z-value**	**Improvement when temporal derivative added**	**Improvement when dispersion derivative added**
Alpha-Chloralose	Very large	Overshoot too small	Delayed	Too wide	Medium	Small	Large
Pentobarbital	Small	Good	Very delayed	Slightly too wide	Low	Large	Small
Ketamine	Medium	Overshoot slightly too small	Good	Slightly too wide	High	None	Small
Fentanyl	Large	Overshoot slightly too small	Good	Good	Very high	None	Very small
Isoflurane	Large	Overshoot too small	Slightly too slow	Too wide	Medium	Very small	Large
Propofol	Very small	Good	Delayed	Too wide	Low	Medium	Medium

Table 2.1: Descriptions of the observed mean modulatory effects of the anaesthetics in the anaesthetics dataset on hemodynamics (compared to the canonical [ΔHbR]) and their expected consequences for GLM Z-values (GLM Z-value is positively correlated with response size and response fit, where response fit depends on how similar the response shape, onset and width are compared to canonical [ΔHbR]).

Anaesthetic	**Z-value using canonical [ΔHbR]**			
		+ temporal derivative	**+ dispersion derivative**	**+ both derivatives**
Alpha-Chloralose	4.9246***	4.7261***	6.0561***	5.7941***
Pentobarbital	-1.0209	-1.2186	-0.0108	-0.3192
Ketamine	3.6614***	3.6956***	3.8515***	3.7726***
Fentanyl	5.1256***	4.8869***	5.1434***	5.1062***
Isoflurane	6.1086***	6.1642***	6.4238***	6.3855***
Propofol	2.5422**	2.4240**	3.1038**	2.9854**

Table 2.2: Multi-subject neural activation Z-values under each anaesthetic using GLM with canonical [ΔHbR] and different combinations of its temporal and dispersion derivatives as basis functions. * equivalent to $p < 0.05$, ** equivalent to $p < 0.01$ and *** equivalent to $p < 0.001$.

Anaesthetic	**Percentage of activations detected as significant value using canonical [ΔHbR]**			
		+ temporal derivative	**+ dispersion derivative**	**+ both derivatives**
Alpha-Chloralose	82%	82%	100%	91%
Pentobarbital	0%	0%	20%	20%
Ketamine	67%	67%	89%	78%
Fentanyl	100%	100%	100%	100%
Isoflurane	100%	100%	100%	100%
Propofol	58%	50%	67%	67%

Table 2.3: Percentage of rat brain hemispheres detected to have significant (equivalent to $p < 0.05$) neural activation under each anaesthetic using GLM with canonical [ΔHbR] and different combinations of its temporal and dispersion derivatives as basis functions.

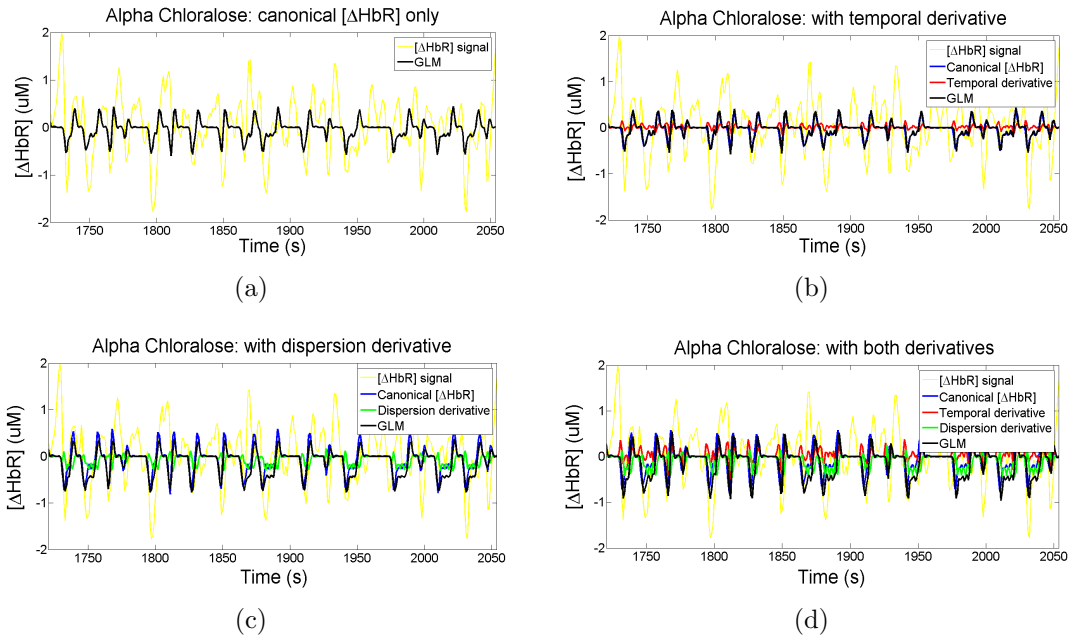

Figure 2.7: Example GLM fits for Alpha-Chloralose when different basis functions are included in the GLM. The Z-values for this example are (a) 13.1909***, (b) 10.8094***, (c) 14.7342*** and (d) 13.9320***.

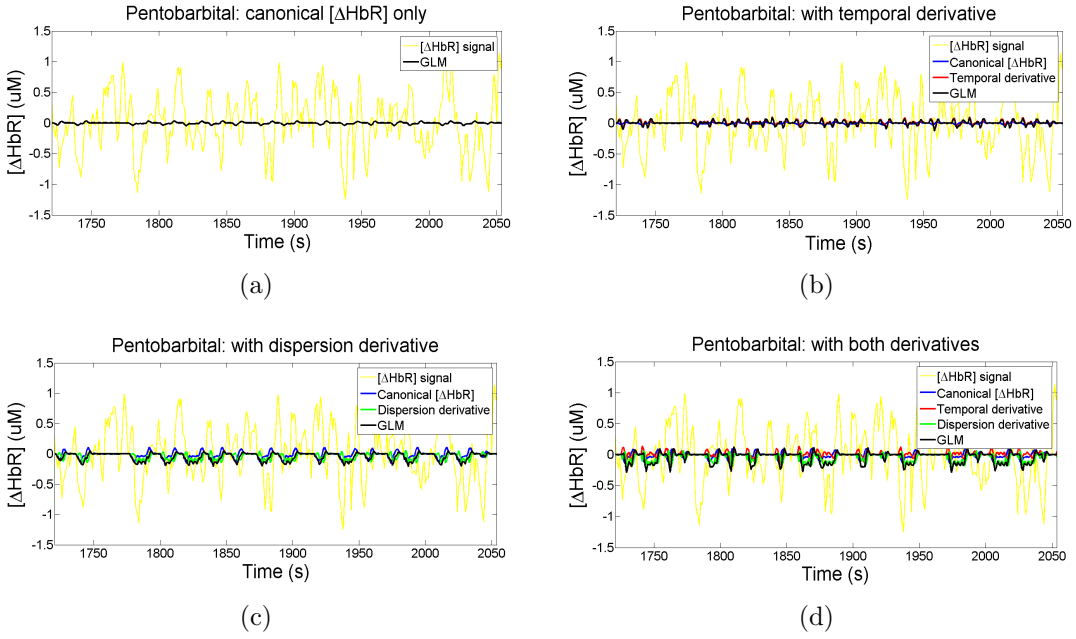

Figure 2.8: Example GLM fits for Pentobarbital when different basis functions are included in the GLM. The Z-values for this example are (a) 1.1446, (b) 1.2180, (c) 2.3595** and (d) 2.0911*.

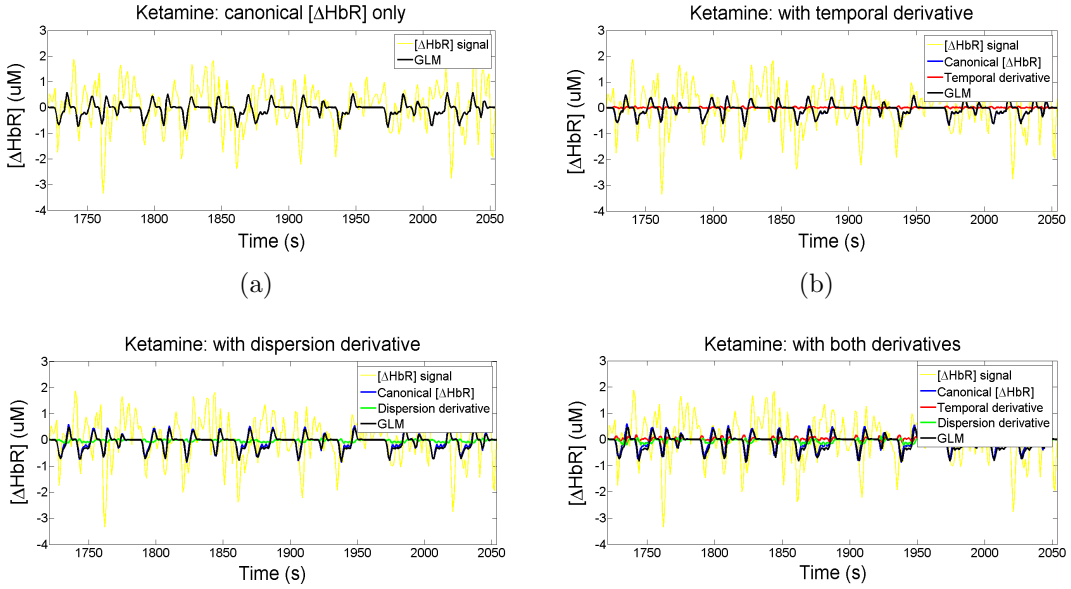

Figure 2.9: Example GLM fits for Ketamine when different basis functions are included in the GLM. The Z-values for this example are (a) 15.6312***, (b) 11.4665***, (c) 12.3551*** and (d) 12.6577***.

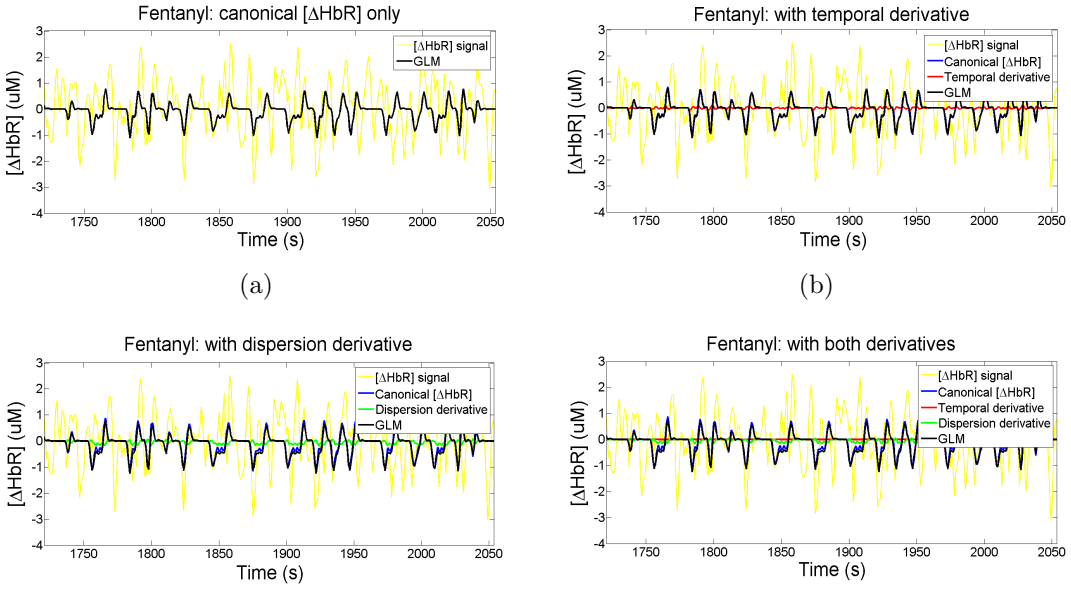

Figure 2.10: Example GLM fits for Fentanyl when different basis functions are included in the GLM. The Z-values for this example are (a) 11.6132***, (b) 11.7937***, (c) 12.2008*** and (d) 12.2093***.

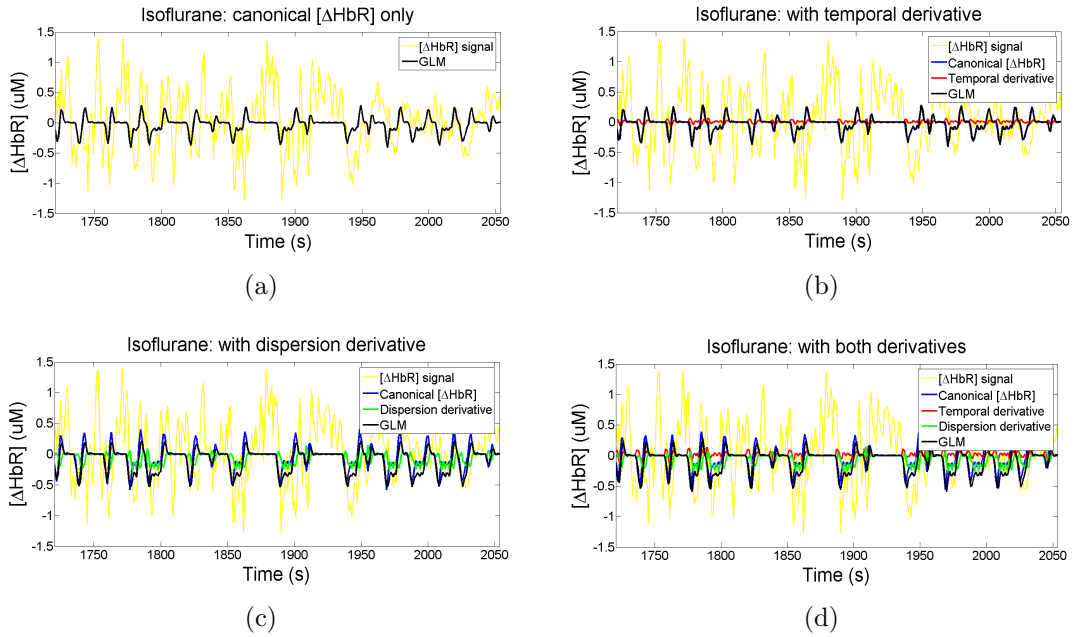

Figure 2.11: Example GLM fits for Isoflurane when different basis functions are included in the GLM. The Z-values for this example are (a) 9.3979***, (b) 9.0702***, (c) 11.3190*** and (d) 10.2545***.

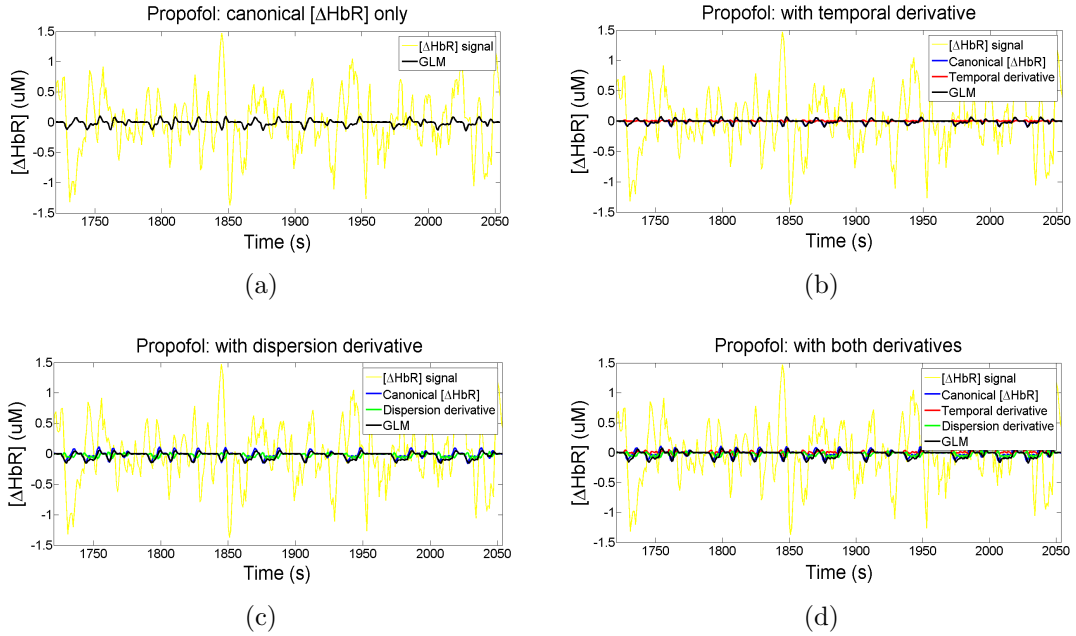

Figure 2.12: Example GLM fits for Propofol when different basis functions are included in the GLM. The Z-values for this example are (a) 4.3866***, (b) 2.6067**, (c) 2.4794** and (d) 3.3017***.

- There is a small improvement to the Z-value under Isoflurane when the temporal derivative is added.

However, the results differ from the qualitative observations in that:

- Ketamine has a lower than expected Z-value when only canonical [ΔHbR] is included as a basis function in the GLM. This is because, even though the average GLM β value under this anaesthetic is larger than the average value under any of the other anaesthetics (9.3220, see Table 2.4), the variance of the individual $\boldsymbol{\beta}$ is also largest under this anaesthetic (standard deviation of 42.6347). This causes the group-level Z-value to be smaller than would be expected from the mean response (see Fig.2.6) because group-level Z-values are a function of the variance of the individual Z-values as well as their mean. This large variance indicates that the actions of Ketamine on hemodynamics are more individualised than the other anaesthetics.
- The predicted group-level Z-values under Isoflurane and Propofol were underestimates rather than overestimates. This is because the variances of the individual $\boldsymbol{\beta}$ under these anaesthetics are small (standard deviations of 5.6276 and 4.8396, respectively, see Table 2.4), i.e. there was little variation in the hemodynamic responses generated across the rats under these two anaesthetics.
- Adding the temporal derivative lowered the Z-values for all anaesthetics except for Ketamine and Isoflurane, which experienced very small improvements to their Z-values. Despite the prediction that adding derivatives would only improve Z-values because they ought to reduce the standard error of any individual model fit (increasing the individual $\boldsymbol{\beta}$ and therefore increasing their average value), there are three ways in which a derivative can reduce a group-level Z-value: 1) by increasing the variance of the individual $\boldsymbol{\beta}$ because the heterogeneity in the response onset/width can

now be fitted to more accurately (meaning that the response amplitude can be fitted to more accurately too), 2) by removing the overestimation of Z-values resulting from mismatches in response onset/width being modelled as additional response amplitude[1], 3) by worsening the accuracy of the response fit as the flexibility endowed by the derivatives allows for noise to be fitted to to a greater extent. As Table 2.4 presents a complicated picture of the effects of the temporal derivative on the mean and variance of the individual $\boldsymbol{\beta}$ (with some anaesthetics experiencing increases or decreases in both the mean and variance or an increase in one and a decrease in the other), it would appear that temporal derivatives alter GLM group-level Z-values in a mixture of ways across the six anaesthetics.

The overall learning to be made from these results is that the prediction that adding derivative basis functions to a GLM always increases Z-values is false. At this stage in the generative embedding algorithm spanning this and the next chapter, no conclusions can be made about the accuracies of the Z-values listed in Table 2.2, and so it remains to be seen which, if any, derivative(s) improve the accuracy of GLM Z-values under the various anaesthetics most.

The final and most important observation to be made from this chapter is that there is considerable variation in the Z-values across the six anaesthetics, with some anaesthetics causing the hemodynamics responses from all subjects to go undetected (Pentobarbital, see Table 2.3), whereas others (Fentanyl and Isoflurane) are consistently detected as significant regardless of whether temporal and/or dispersion derivatives are included as basis functions in the GLM. This variation across anaesthetics may be caused by:

[1]By simulating BOLD responses with ground-truth differences in response amplitude (up to 50%), time-to-onset (up to 3s) and width (up to 4s), Lindquist and Wagner, 2007, calculated the confusability (bias in response amplitude induced by true changes in the response time-to-onset and width) suffered by the standard GLM. They found that 95% of the response delay was detectable from the response amplitude (i.e. captured by the canonical HRF) and 5% of the amplitude and 100% of the width was detectable from the delay (i.e. captured by the temporal derivative). Only 12% of the response width was detected by the dispersion derivative.

Anaesthetic	**Canonical [ΔHbR]**							
	μ	σ	**+ temporal derivative** μ	**+ temporal derivative** σ	**+ dispersion derivative** μ	**+ dispersion derivative** σ	**+ both derivatives** μ	**+ both derivatives** σ
Alpha-Chloralose	6.2028	16.8644	5.8156	16.2953	7.9688	17.1699	7.5998	17.7761
Pentobarbital	-0.4523	3.5105	-0.8148	3.4274	-0.0312	3.6604	-0.3192	3.6328
Ketamine	9.3220	42.6347	7.1602	26.8794	8.0133	29.7856	7.9193	31.4552
Fentanyl	6.6565	10.3522	6.5074	11.1355	6.8756	10.0101	7.0386	12.2589
Isoflurane	5.7716	5.6276	6.1472	6.7532	7.3062	6.5279	6.8789	6.0685
Propofol	1.6666	4.8396	1.6307	5.0411	2.6020	4.1976	2.4568	3.8179

Table 2.4: Mean, μ, and standard deviation, σ, of individual neural activation Z-values under each anaesthetic for GLM with canonical [ΔHbR] and different combinations of its temporal and dispersion derivatives as basis functions.

1. Anaesthetics altering neuronal activity (the interest variable) to different extents,
2. Anaesthetics altering neurovascular coupling to different extents, confounding GLM analyses,
3. Anaesthetics altering both neuronal activity and neurovascular coupling to different extents, confounding GLM analyses.

Given the evidence in the literature that anaesthetics act differently on vascular processes as well as neuronal processes (see Table 8.1), it is hypothesised that differences in neurovascular coupling are at least partly responsible for the different GLM significances across anaesthetics. This hypothesis will be tested in Chapter 3.

Chapter 3

Testing BOLD-fMRI GLM sensitivity when accounting for anaesthetic neuromodulation of neurovascular coupling

Conclusions about neuronal activity made from fMRI studies are usually inferred from GLM Z-value significances only. It is therefore critical that these Z-value significances reflect neuronal activity accurately. In Chapter 2 it was found that the significance with which neuronal activity was detected using GLM depended on the anaesthetic neuromodulation under which the hemodynamics were recorded. In this chapter, the aim is to investigate whether these differences in significances across the anaesthetics reflect differences in neuronal activity only, as is often assumed to be the case, or if they also/instead reflect differences in neurovascular coupling. This will be achieved by simulating "unanaesthetised" hemodynamics responses to anaesthetic-specific neuronal activities and comparing the Z-value significances with which GLM detects the neuronal activity driving these responses under each anaesthetic with the Z-value significances it detected

for the real responses in Chapter 2. The implications that anaesthetic neuromodulations have for the interpretation of GLM results should then be clear.

3.1 Generative modelling

As in Brodersen et al., 2010, and Brodersen et al., 2011, Dynamic Causal Models (DCMs) were the generative models of choice for this generative embedding algorithm. This is not only because they are the most widely used generative models in neuroimaging studies, but also because they have multimodal variants, meaning that generative models are available for fitting to both the neuronal (DCM-EEG) [41] and hemodynamic (DCM-fMRI) [59] signals within a single framework [18][19].

3.1.1 Single-modality DCMs

In neuroimaging research, DCMs are typically used to infer brain connectivity from signals originating from distinct regions of the brain, with some comparison being made about the strengths of these connections in pathology (example studies include [118] and [13]). However, as the ultimate goal of this work is to understand the effects of neuromodulations on neurovascular coupling, it is important that the effects of neuromodulations on brain connectivity are ignored and instead their effects on the local biophysical processes are focussed on. For this reason, it is sufficient for this work that only the EEG and [ΔHbR] signals recorded from a single brain region (as is available in the anaesthetics dataset) are analysed.

There are two families of DCM: those related to neural electrophysiology and those related to neural hemodynamics, as shown in Fig.3.1. They all have the same two-level structure (which together make the complete forward model):

1. The neural level, where the dynamics of the neuronal activity within a cortical column are modelled as:

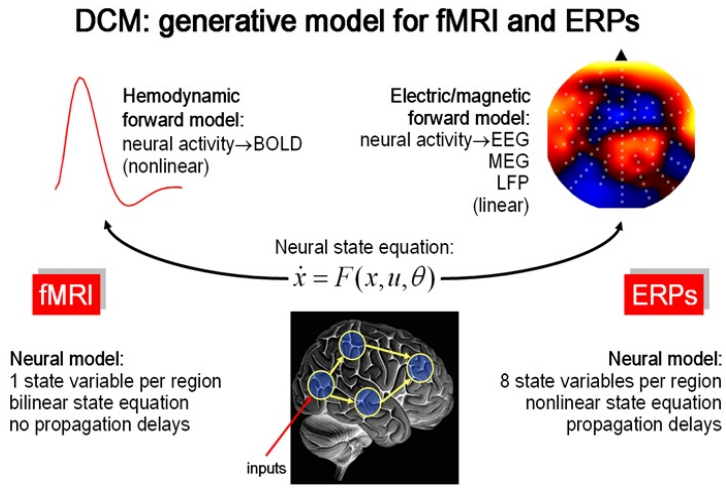

Figure 3.1: The two families of DCM: those fitted to measures of neural hemodynamics (left branch), and those fitted to neural electrophysiology (right branch). Figure reproduced from SPM8 manual, Figure 32.1 [6].

$$\dot{x} = f(\boldsymbol{x}, u, \boldsymbol{\theta_N}), \tag{3.1}$$

where $\boldsymbol{x}$ are the mean transmembrane potentials and currents of neuronal areas, u is the known external input driving the neuronal activity and $\boldsymbol{\theta_N}$ are biophysical parameters.

Because electrophysiology signals such as EEG are direct measures of neuronal activity, this level of their DCMs are complex, as subtleties in their neural model give rise to distinct features in the resulting imaging signal. The neural model of DCM-EEG, which is fitted to the EEG signals in this work, is based on the Jansen-Rit neural mass model [84]. This model was chosen because its ability to produce outputs similar to both spontaneous EEG α oscillations and evoked potentials had been extensively established in the literature. This neural model models the electrical dynamics of a cortical column by grouping the different cell populations within the column into three cell populations: excitatory interneurons termed spiny stellate cells corresponding to the granular layer 4, and inhibitory interneurons and excitatory pyramidal cells corresponding to the supra- and infra-granular

layers respectively.

External input, u, is relayed to the spiny stellate cells via the thalamus, resulting in a perturbation of their average membrane potential, v, that is modelled by the univariate higher-order differential equation (see Fig.3.2 for parameter descriptions):

$$\ddot{v} = \frac{H_e}{\tau_e}u(t) - \frac{2}{\tau_e}\dot{v} - \frac{1}{\tau_e^2}v. \tag{3.2}$$

Equ.3.2 can be represented as the multivariate first-order set of differential equations:

$$\dot{v} = c, \tag{3.3}$$

$$\dot{c} = \frac{H_e}{\tau_e}u(t) - \frac{2}{\tau_e}c - \frac{1}{\tau_e^2}v, \tag{3.4}$$

where c is proportional to the current across the spiny stellate cells.

This is the form of the equations used in DCM-EEG to model the dynamics of the membrane potentials of all three cell populations.

Because all of the cell populations are connected to each other, the spiny stellate cells perturb the membrane potentials of the inhibitory interneurons and excitatory pyramidal cells. The dynamics of these biophysically realistic internal electric connections depend on two operators: $S(x)$ and $P(t)$. $S(x)$, which is a sigmoidal function, transforms the potential, x, of a firing cell population into a firing rate, its slope reflecting the synchronicity of the firing, whereas $P(t)$, which is an impulse response, transforms the average density of these presynaptic inputs into the average postsynaptic membrane potential of the receiving cell population.

The average membrane potential of all three cell populations take the same form as Equ.3.2, with inhibitory parameters replacing the excitatory

parameters where necessary.

For the more common uses of DCM (to make inferences about brain connectivity), the brain region(s) directly perturbed by the external input can themselves perturb other brain regions. However, because only a single-region DCM is of interest to this work (as the SEPs are recorded from a single location only), these connections do not have to be incuded in the DCM-EEG used in this thesis.

Signals of hemodynamics are considerably further removed from the neuronal activity than neuronal signals. For this reason, single-source DCM-fMRI has only a linear neural model describing the dynamics of one cell population, avoiding the model being unduly complex.

2. The observational level, which converts the electrical activity, x, into the signal recorded, h:

$$h = g(\boldsymbol{x}, \boldsymbol{\theta_O}) + N, \tag{3.5}$$

where g is the observation function (a system of ordinary differential equations), $\boldsymbol{\theta_O}$ are more biophysical parameters and N is measurement noise. For DCM-EEG, $\boldsymbol{\theta_O}$ is simply a vector of gain factors representing the attenuation of the electrical currents from the three cell populations through the skull. The full (single-region) DCM-EEG specific to EEG responses to somatosensory stimulation is shown in Fig. 3.2 along with biological interpretations of its neural model parameters, $\boldsymbol{\theta_N}$, and observer model parameters, $\boldsymbol{\theta_O}$ [41].

In early versions of SPM, only the transmembrane potential of the pyramidal cells were modelled to give rise to the EEG signal as they are the most superficial cell population and their apical dendrites have tangential arrangement to the cortical surface (so that their dipolar fields reinforce rather than cancel each other out). There is a growing body of evidence

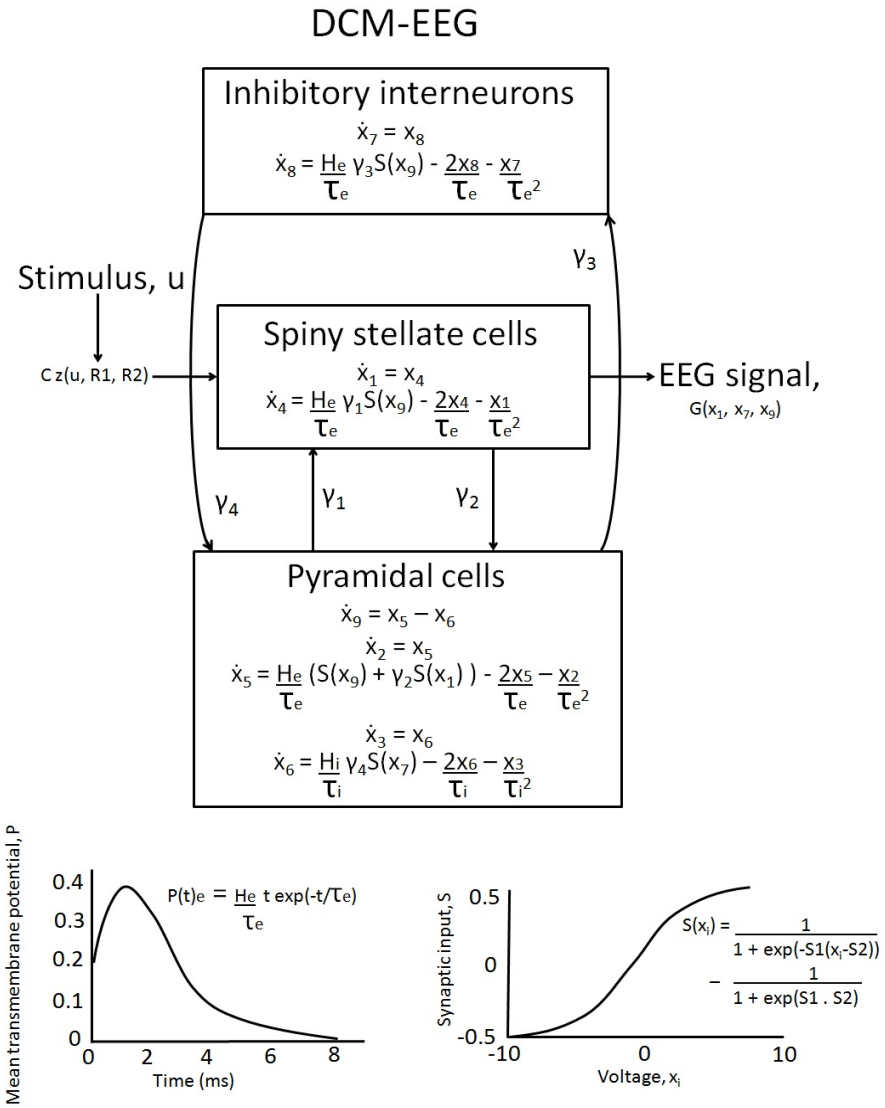

Figure 3.2: Single region DCM-EEG for SEPs driven by exogenous input u. Operators $S(x)$, which transforms the potential, x, of a firing cell population into a firing rate, and $P(t)$, which transforms the average density of the presynaptic inputs into the average postsynaptic membrane potential of the receiving cell population over time t are also shown. z is a Gaussian function that disperses u. Neural model free parameters, $\theta_{\mathbf{N}}$: H_e = excitatory synaptic density; H_i = inhibitory synaptic density; $\gamma 1$ = excitatory to excitatory connection strength; $\gamma 2$ = excitatory to excitatory connection strength; $\gamma 3$ = excitatory to inhibitory connection strength; $\gamma 4$ = inhibitory to excitatory connection strength; C = modulation of input strength; R1 = onset of stimulus input; R2 = dispersion of stimulus input; S1 = activation function threshold; S2 = activation function dispersion; τ_e = excitatory synaptic time constant; τ_i = inhibitory synaptic time constant. Observer model free parameters, $\boldsymbol{\theta_O}$: $\mathbf{G}$ = vector of gain parameters quantifying the attenuation of the electrical currents through the skull.

suggesting that the other cell populations also contribute to the EEG signal, even though they display more closed-field dendritic configurations [122][27]. For this reason, the EEG signal in SPM8 is modelled as a linear weighting of all three cell populations. Because there is little consensus about the relative contributions of these cell populations to the EEG signal, their linear weightings are free parameters in DCM-EEG, with prior means of 0.6 for pyramidal cells, 0.2 for spiny stellate cells and 0.2 for inhibitory interneurons[1] but no prior second moments. As the observational model of DCM-EEG does not feature in Combined DCM, these weighting parameters are of no interest to this work. Therefore, to prevent these little-understood parameters from influencing the Combined DCM fits, it was decided to fix them to their prior values.

For DCM-fMRI, the observational level incorporates the extended Balloon Model [25], which models veins as expandable, balloon-like compartments that swell when brain activation occurs. This increase in venous cerebral blood volume (CBV) is caused by the pressure gradients that form across the veins in response to increased cerebral blood flow (CBF), which in turn is the response to oxygen metabolism (fuelling the neuronal activity) in the capillaries. These CBF and CBV increases result in an oversupply of oxygenated blood (increased CBF means that there is less time for oxygen to be extracted by the neighbouring tissue), leading to an overall decrease in HbR concentration.

The default DCM-fMRIs fitted using SPM8 have only three free parameters at this observational level: decay (decay of vasodilatory signal from neuronal activity), transit (transit time of blood in veins) and epsilon (ratio of intra- and extra-vascular components of the gradient echo fMRI signal). As epsilon is not relevant to HbR signals, only two parameters were

[1] Dr. Rosalyn Moran explained in a personal correspondance that these prior values were chosen because they optimised DCM-EEG fits to real data.

informative for this work, limiting the possible biological insights that could be made. Therefore, the code was adapted to estimate more free parameters, as was done in earlier versions of SPM. The SPM8 code was also adapted to fit DCMs to [ΔHbR] data (in Molars = moles/L) rather than BOLD signal:

$$\Delta(\frac{Q}{V}) = \frac{156E0}{66,500}(\frac{q}{v} - 1), \tag{3.6}$$

where $\Delta\frac{Q}{V}$ is the change in HbR concentration (Q is HbR content and V is venous blood volume fraction in veins), $E0$ is the resting oxygen extraction fraction from the blood in the capillaries, and q and v are the normalised HbR content and volume fraction (relative to resting), respectively. This expression was derived by noticing that:

$$\Delta(\frac{Q}{V}) = (\frac{Q}{V}) - (\frac{Q_0}{V_0}) = \frac{Q_0}{V_0}(\frac{q}{v} - 1), \tag{3.7}$$

where Q_0 and V_0 are the HbR content and venous blood volume at rest, and $Q_0 = V_0(1 - Y_0)$[Hb], where Y_0 is the oxygen saturation of the blood at rest (Y_0 = 1 - E0) and [Hb] is the effective total hemoglobin concentration in blood, which is estimated to be $\frac{156}{66,500}$ because there is about 156g/l, of hemoglobin in blood, and 66,500g/mole of hemoglobin in adult rats, giving $\frac{156}{66,500}$ moles/l [128][49][154].

To highlight that a number of modifications were made to DCM-fMRI in its implementation within Combined DCM, this model, which is shown in Fig.3.3 along with descriptions of its free parameters, is referred to as modified DCM-fMRI (mDCM-fMRI) [59].

3.1.2 Combining DCM-EEG and mDCM-fMRI

In DCM, the probability densities of the parameterised neuronal activity and hemodynamics biological processes are approximated by multivariate Gaussians

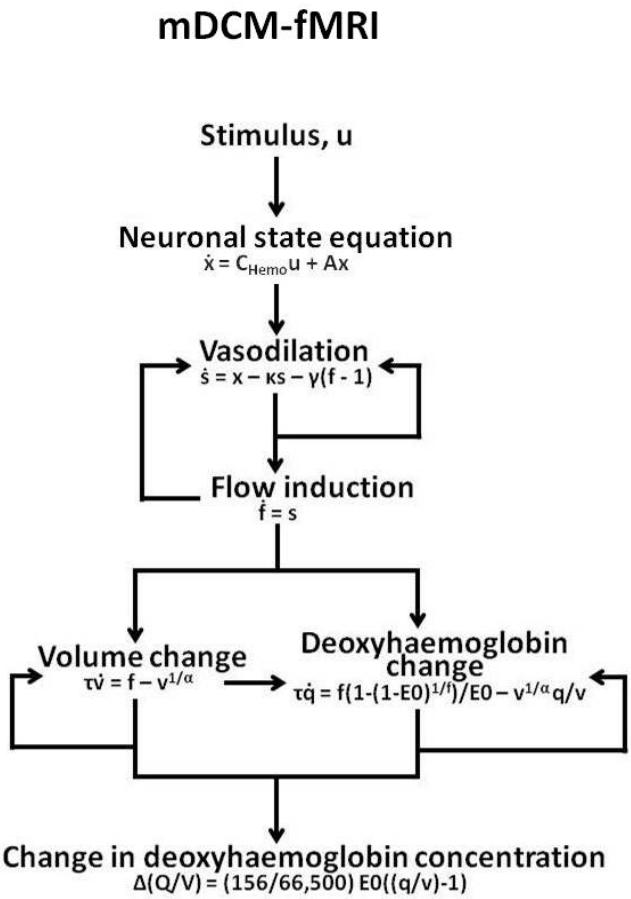

Figure 3.3: Single region mDCM-fMRI for [ΔHbR] signals. Neural model free parameters, $\boldsymbol{\theta_N}$: A = self-inhibition of the electrical activity, x, and C_{Hemo} = amplification of the effect of the exogenous input u on x. Observer model free parameters, $\boldsymbol{\theta_O}$: κ = rate of vasodilatory signal decay; γ = rate of bloodflow autoregulatory feedback; τ = blood transit time in venous compartment; α = Grubbs vessel stiffness exponent; E0 = resting oxygen extraction fraction.

The differential equations for q and v derive from mass balance considerations: the balloon will swell (v will increase) until the blood flow out of the balloon ($v^{\frac{1}{\alpha}}$) matches the blood flow in (f) (steady state), and q reflects the difference in the rates of HbR entry (blood flow in, f, multiplied by the proportion of oxygen extracted from the blood before the balloon, i.e. through the capillaries) and clearance (blood flow out, ($v^{\frac{1}{\alpha}}$), multiplied by average HbR concentration in the balloon, $\frac{q}{v}$).

(see Appendix B.2 for details about SPM's Variational Bayes algorithm used for DCM fitting - the basic structure of this algorithm is adopted for fitting Combined DCMs later in this chapter). Therefore, the effects of neuromodulations on neurobiology may be captured by the following moments:

- Expectations of neuronal activity parameters, $\mathbf{Ep}_{EEG}$
- Covariances of neuronal activity parameters, $\mathbf{Cp}_{EEG}$
- Expectations of hemodynamics parameters, $\mathbf{Ep}_{fMRI}$
- Covariances of hemodynamics parameters, $\mathbf{Cp}_{fMRI}$
- Covariances between the neuronal activity and hemodynamics parameters, $\mathbf{Cp}_{EEG-fMRI}$

As already explained, mDCM-fMRI comprises a neural model as well as an observer model. However, the simplicity of the mDCM-fMRI neural model (self-inhibitory differentiation of the exogenous input - see Fig.3.3), renders it inadequate for capturing the anaesthetic-specific effects on neuronal activity[1]. For this reason, it is necessary to fit a more detailed neural model, such as DCM-EEG, to the EEG data. $\mathbf{Cp}_{EEG-fMRI}$ cannot be estimated from the separate fitting of DCM-EEG to the EEG data and mDCM-fMRI to the [ΔHbR] data, and so, in order to estimate anaesthetic-specific values for $\mathbf{Cp}_{EEG-fMRI}$, a Combined DCM needs be constructed.

There is experimental evidence that a complex mix of mechanisms mediate the coupling of neuronal activity and hemodynamics [8], giving rise to a number of linking models in the literature. For example, Riera et al.'s 2006 model uses concentration of Nitric Oxide generated by all of the neuronal cell populations to

[1]The neural model of DCM-fMRI is simple because it is optimised for calculating connection strengths between brain regions from hemodynamics responses only (which are a considerably filtered version of the underlying neuronal activity) and not for accurately modelling the biophysical processes in each brain region. However, because the anaesthetics dataset comprises concurrent measures of the neuronal activity and hemodynamics, it is appropriate for this work to fit a Combined DCM with a more complex neural model than in mDCM-fMRI.

drive CBF [162], whereas Sotero et al.'s 2005 model uses the oxygen and glucose metabolism resulting from excitatory neuronal activity [86]. This lack of consensus about which aspects of neuronal activity drive evoked hemodynamics motivated a search for a model linking DCM-EEG to mDCM-fMRI that generated adequate fits to the [ΔHbR] signals whilst minimising the number of extra parameters needed for its implementation, as recommended by Occam's Razor. This search, which culminated in Sotero et al.'s metabolism model being chosen to link DCM-EEG and mDCM-fMRI in Combined DCM, is detailed in Appendix C.

3.1.2.1 Details of Combined DCM

In Sotero's metabolism model, different aspects of the hemodynamics are driven by excitatory, U_e, and inhibitory, U_i, neuronal activities, as shown in Fig.3.4, with U_e and U_i corresponding to the average numbers of synapses that will receive an action potential in one time unit, i.e. synaptic activity:

$$U_e = (\gamma_1 + \gamma_3)S(x_9) + \gamma_2 S(x_1); \tag{3.8}$$

$$U_i = \gamma_4 S(x_7); \tag{3.9}$$

where $\gamma_i (i = 1, ..., 4)$ = total numbers of synapses between the cell populations (strengths of the intrinsic connections) and $S(x_i, i = 1, ..., 9)$ = cell population firing rates in DCM-EEG.

Driving hemodynamics using synaptic activity agrees with the findings by Logothetis et al., 2001, that BOLD signal more strongly reflects neuronal input and processing than neuronal output: convolution of LFP signals, which reflect synaptic activity, with an impulse response function, generated more accurate hemodynamics than when multi-unit activity signals, which reflect spiking, were convolved instead [104]. However, experiments carried out by Heeger et al., 2000,

and Smith et al., 2002, supported the opposite conclusion [74][178]. To gain some understanding of these conflictory observations, Rosa et al., 2011, used modelling approaches similar to those used in this work to investigate the circumstances in which synaptic and spiking activities contribute to the BOLD signal. Specifically, they used the Variational Bayes algorithm designed to fit single-modality DCMs to neuroimaging data (described in Appendix B.2) to fit three different neurovascular coupling models to human EEG-BOLD data:

1. Synaptic input model: CBF is driven by the concentration of Nitric Oxide generated by the neuronal cell populations, [162]
2. Spiking output model: CBF is driven by firing rate of pyramidal cells,
3. Mixed model: CBF is driven by both concentration of Nitric Oxide and pyramidal cell firing rate.

They then used Bayesian model selection to conclude that the BOLD signal is most dependent on synaptic activity during low frequency neuronal activity.

The different influences of U_e and U_i on the hemodynamics are experimentally motivated. Firstly, Sotero et al.'s model assumes that CBF is not coupled to inhibitory neuronal activity. Evidence supporting this assumption includes the findings of [97] and [114] that show that GABA, the major inhibitory neurotransmitter, is not involved in activation-induced CBF increases, whereas some of the products of glutamate (the major excitatory neurotransmitter) recycling are vasoactive (see Fig.1.2). Secondly, Sotero et al.'s model assumes that no glycogenolysis pathway exists to fuel removal of GABA from the synaptic cleft, whereas such a pathway does exist to fuel glutamate clearance (part of the glutamate-glutamine cycle), explaining the differences between the equations describing oxygen consumption for excitatory and inhibitory activities, $m_e(t)$ and $m_i(t)$, in Fig.3.4. This assumption was made based on experiments showing that glycolysis (catabolism of glucose), could not explain lactate accumulation during

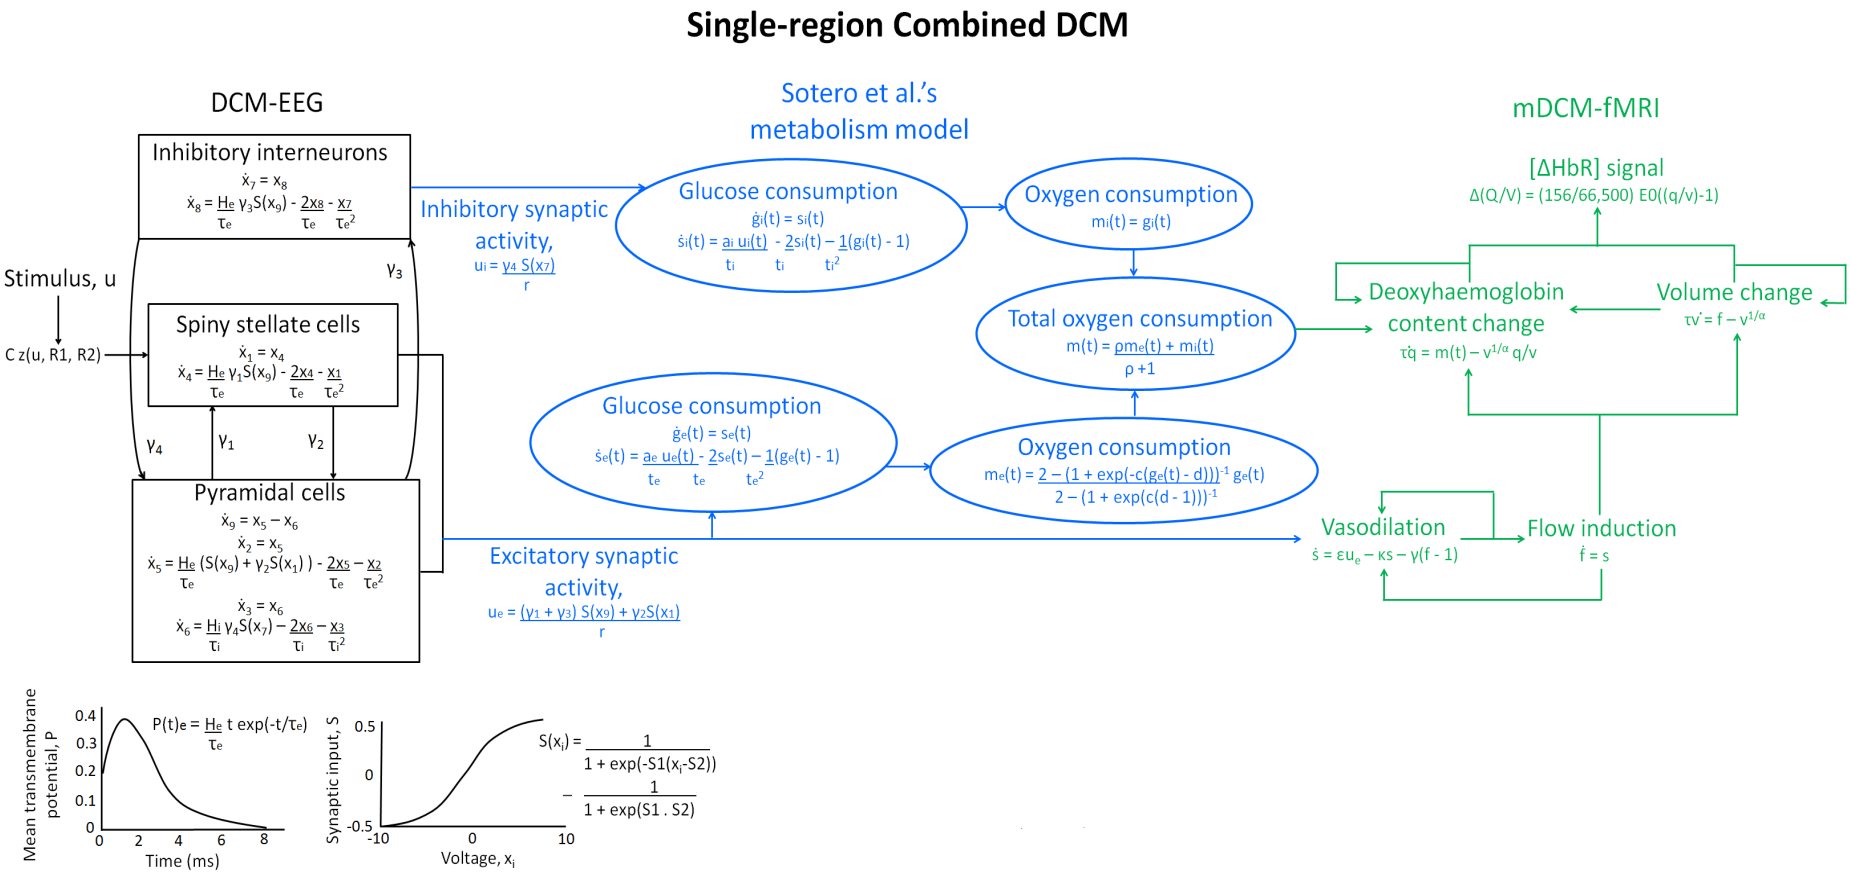

Figure 3.4: Single-region Combined DCM, which uses a modified version of Sotero et al.'s metabolism model to link DCM-EEG to mDCM-fMRI. Descriptions and prior moment values of the parameters are listed in Table 3.1.

sensory stimulation. See Fig.3.5(a) for an explanation of how glycolysis creates [155]:

1. ATP to fuel glutamate-glutamine cycling,
2. Lactate to fuel other aspects of neuronal activity[1].

It is therefore hypothesised that another pathway exists for generating ATP quickly during intense neuronal activation, namely glycogenolysis (catabolism of glycogen, a polymer of glucose), as described in Fig.3.5(b) [175].

Because it has been shown that glutamatergic neurons outnumber their GABAergic counterparts ten to one in human cortex [2], Sotero et al.'s model does not include this Glycogen Shunt model for fuelling inhibitory neuronal activity. Therefore, the oxygen-to-glucose index for GABAergic neuronal activity is $OGI|_i = 6$, owing to six oxygen molecules being necessary to oxidise one glucose molecule ($C_6H_{12}O_6$):

$$C_6H_{12}O_6 + 6O_2 \rightarrow 6CO_2 + 6H_2O. \tag{3.10}$$

In contrast, the oxygen-to-glucose index for glutamatergic neuronal activity is $OGI|_e = 6 - 3x$, where x = fraction of glucose created in the Glyogen Shunt rather than undergoing glycolysis. The fraction of glucose involved in the Glycogen Shunt is modelled as a sigmoidal function of the glucose consumed for excitatory activity, g_e, in order to capture the increasing usage of the glycogenolitic pathway with increasing excitatory neuronal activity:

$$x = \frac{1}{1 + e^{-c(g_e(t)-d)}}, \tag{3.11}$$

where c= steepness and d = position of sigmoid.

[1]The majority of ATP is used to restore the ions fluxes underlying synaptic potentials and action potentials via ion pumps.

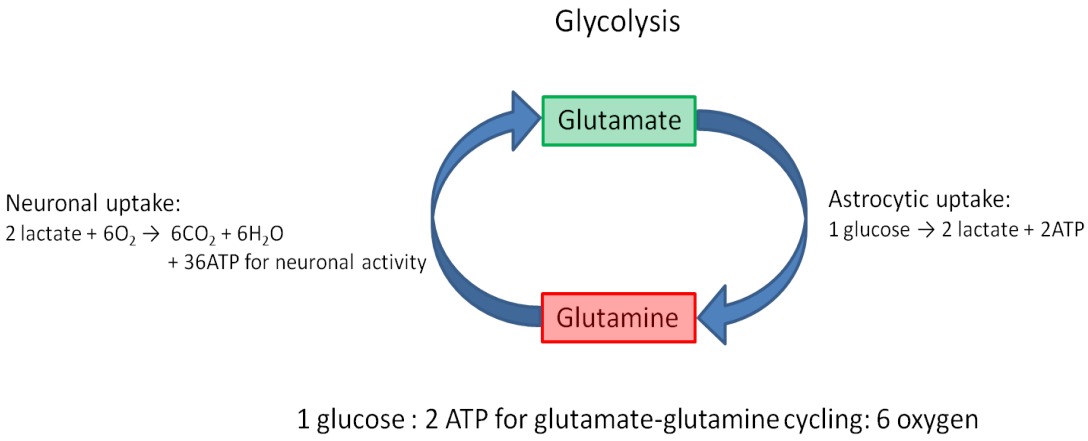

(a)

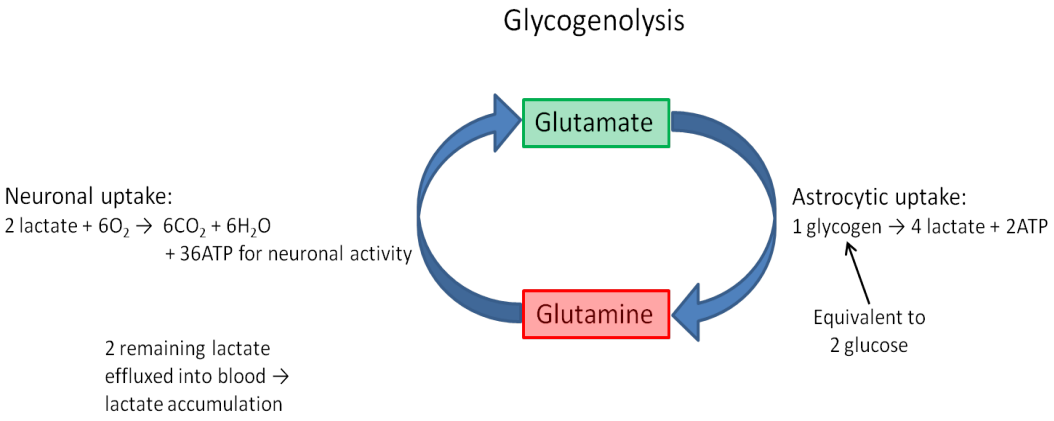

(b)

Figure 3.5: In the Glycogen Shunt model, glutamate-glutamine cycling provides energy for neuronal activity via two pathways: (a) glycolysis and (b) glycogenolysis. Glycogen is a polymer of glucose molecules, and so glycogen comprising two glucose molecules is required to create the two moles of ATP necessary for the processes occurring in the astrocytes. If fraction x of glucose is involved in glycogenolysis then one glucose molecule generates $(2 - x)$ ATP molecules. As three oxygen molecules are metabolised for every ATP molecule in the oxidisation reaction, this results in $OGI|_e = 6 - 3x$.

3.1.2.2 Modifications to Sotero et al. linking model

One way that the linking model of Combined DCM differs from the model proposed by Sotero et al. is that it includes a parameter to scale both the excitatory and inhibitory neuronal activities. This modification was necessary because, in their work, Sotero et al. simulated normalised neuronal activities, u_e and u_i, i.e. neuronal activities relative to baseline levels. Rather than change the prior values for all of the original parameters in Sotero et al.'s model in order to fit to the unnormalised EEG signals in the anaesthetics dataset, it was instead decided that one more parameter, r, would be added to Combined DCM to scale U_e and U_i (see Fig.3.4). Separate parameters for scaling U_e and U_i were not introduced in order to reduce the number of (potentially very redundant) parameters in Combined DCM.

Scaling U_e and U_i by parameter r in DCM-EEG is not equivalent to normalising the U_e and U_i generated in Sotero et al.'s model because of the difference in sigmoidal functions used to transform transmembrane potential to firing rate in the two models: $S(x)$ in DCM-EEG models negative firing rates (relative to baseline levels) generated from cells with negative transmembrane potentials, whereas it is constrained to (non-biophysically) only take positive values in Sotero et al.'s model. By passing scaled rather than normalised synaptic activities to Sotero's metabolism and the extended balloon model, it was unnecessary to subtract 1 from their influences on glucose metabolism and the vasodilatory signal, as is done in Sotero et al.'s model. This is the second and final way in which the linking model of Combined DCM differs from Sotero et al.'s model.

3.1.2.3 Combined DCM parameters

The prior expectation values of the DCM-EEG parameters used by default in SPM8 are based on a combination of experimental evidence reported in the literature (recorded from a range of animals) and analytical observations of model

dynamics made by Jansen and Rit, 1995 [84]. As no consistent conclusions have been made about how the biological processes modelled by neuronal mass models vary across species, these prior expectation values were adopted for the Combined DCMs fitted to the rat hemodynamics responses in this chapter.

It would seem most relevant for the posterior parameters from Combined DCM fitting to the awake-rat canonical [ΔHbR] to be used as priors in Combined DCMs fitted to rat hemodynamics responses in this chapter because the prior expectation values of the DCM-fMRI parameters used by default in SPM8 were optimised analytically for human hemodynamics [58]. Despite this fact, the default prior expectation values for the DCM-fMRI parameters were adopted for this work because they resulted in much more stable fits[1] and reduced the occurrence of nonsensical posterior parameters (such as E0 ¿1) without noticeably compromising the goodness of model fits made using the Variational Bayes algorithm.

In Sotero et al., 2007, the same Volterra kernels used to initially optimise the prior values for DCM-fMRI were used to manually tune the parameters in their metabolism model. To be consistent with the decision to use the prior expectation values of the DCM-fMRI parameters used by default in SPM8, these values have been adopted as the prior expectation values of the parameters of Sotero et al.'s model in Combined DCM too [180].

The only remaining parameter to consider is r, and the prior expectation of this parameter was chosen through trial-and-error fitting of Combined DCM to a sample of the real [ΔHbR] responses. This prior is kept very loose to reflect its large uncertainty. The priors on the 26 other Combined DCM parameters are considerably tighter: the variances of the DCM-EEG and DCM-fMRI parameters used by default in SPM8 are adopted for the corresponding parameters in Combined DCM, and the variances of Sotero et al.'s linking model (with the exception of r, as discussed above) are set to the same value as for the

[1]The stability of DCM-fMRI is explored in Friston et al., 2003 [59].

mDCM-fMRI parameters, justified again by the fact that the same method, namely fitting of Volterra kernels to the same dataset, was used to derive the prior expectation values for these two sets of parameters.

Table 3.1 lists all of the Combined DCM prior moments. Except for S2, all of the parameters are ensured positivity by exponentiation, e.g. $\tau_e = 0.004e^{Ep(\tau_e)}$, where $\mathrm{Ep}(\tau_e) = 0$ in the prior case. This is necessary because the interpretations of these parameters if they took negative values would not be meaningful. The exception is S2 because a negative value of S2 reflects the sigmoid in both the x and y axes, maintaining a positive sigmoidal relationship between transmembrane potential and firing rate.

In DCM-EEG, R1 was not exponentiated, but it was found that fitting Combined DCM would sometimes push R1 to be negative, i.e. for the EEG response to begin before the stimulus. To prevent this from happening, R1 was also exponentiated in Combined DCM.

The profiles of all the DCM-EEG parameters on the SEP response and of all the Combined DCM parameters on the [ΔHbR] response for the range of their posterior values from fitting to the anaesthetics dataset, are shown in Fig.3.6 and Fig.3.7, respectively. It is important to understand that these parameter effects were simulated with the 26 other parameters set to their prior values in order to make the profile of each parameter as clear as possible. As the parameters are inter-dependent, this means that only a very limited picture of the parameter effects can be gained from Fig.3.7, as they can have very different effects on [ΔHbR] signal when the other parameters take different values. What is more, some of the responses are unstable because having the 26 other parameters fixed to their prior values rather than their corresponding posterior values means that an unstable region of the parameter space may be entered. For these plots, zoom-ins are also included to enable easier visualisation of the stable profiles.

It is important to remember that the metabolism model filters the synaptic

Parameter	Interpretation	Model	Prior mean	Prior variance
R1	Onset of stimulus input	DCM-EEG	8	0.0625
R2	Dispersion of stimulus input	DCM-EEG	8	0.0625
C	Modulation of stimulus input strength	DCM-EEG	0.001	0.0313
H_e	Excitatory synaptic density	DCM-EEG	4	0.0625
H_i	Inhibitory synaptic density	DCM-EEG	64	0.0625
τ_e	Excitatory synaptic time constant	DCM-EEG	0.004	0.0625
τ_i	Inhibitory synaptic time constant	DCM-EEG	0.008	0.0625
γ_1	Pyramidal to spiny stellate cells connection strength	DCM-EEG	$\frac{1}{64}$	0.0625
γ_2	Spiny stellate to pyramidal cells connection strength	DCM-EEG	$\frac{1}{64}$	0.0625
γ_3	Pyramidal cells to inhibitory interneurons connection strength	DCM-EEG	$\frac{1}{64}$	0.0625
γ_4	Inhibitory interneurons to pyramidal cells connection strength	DCM-EEG	$\frac{1}{32}$	0.0625
S1	Firing rate threshold	DCM-EEG	1	0.0625
S2	Firing rate dispersion	DCM-EEG	0	0.0625
r	Ratio of neuronal activities during activation to at rest	Sotero	0.01	1000
a_e	Efficiency of glucose consumption response to excitation	Sotero	1.2	0.0183
t_e	Time constant of excitatory glucose consumption	Sotero	1	0.0183
a_i	Efficiency of glucose consumption response to inhibition	Sotero	1.2	0.0183
t_i	Time constant of inhibitory glucose consumption	Sotero	1	0.0183
ρ	Baseline ratio of excitatory and inhibitory synaptic activities	Sotero	5	0.0183
c	Steepness of sigmoid describing glucose involvement in glycogenolysis	Sotero	2.5	0.0183
d	Position of sigmoid describing glucose involvement in glycogenolysis	Sotero	1.6	0.0183
ϵ	Efficacy of bloodflow to excitation	Sotero	0.8	0.0183
κ	Rate of vasodilatory signal decay	mDCM-fMRI	0.64	0.0183
γ	Rate of bloodflow autoregulatory feedback	mDCM-fMRI	0.32	0.0183
τ	Blood transit time through balloon	mDCM-fMRI	2	0.0183
α	Grubb's exponent: blood volume response to bloodflow	mDCM-fMRI	0.32	0.0183
E0	Resting O_2 extraction fraction from HbO	mDCM-fMRI	0.32	0.0183

Table 3.1: List of Combined DCM parameters with their prior moments.

activity simulated by DCM-EEG, and therefore that the DCM-EEG parameters have little influence on the hemodynamics. This is apparent in Fig.3.7, where the DCM-EEG parameters can be seen to only influence the amplitude of [ΔHbR]. It would not be appropriate to fit Combined DCM to [ΔHbR] signals only as the DCM-EEG parameters would be fully redundant, but Combined DCM is suitable when the concurrent EEG signal is also available because the specificities of the DCM-EEG parameters for features in EEG signals are generally high (see Fig.3.6).

3.2 Methods

3.2.1 Combined DCM fitting

The first step of the generative embedding algorithm presented in this chapter is the fitting of Combined DCM to the real anaesthetics data. Because the SNR of the raw signals was insufficient for DCM fitting, and because Combined DCM fitting is computationally intensive, with each fit taking somewhere between three and 18 hours, it was decided that Combined DCM would only be fitted to the mean EEG and [ΔHbR] responses to the 5s stimulation trains recorded from each rat brain hemisphere, i.e. averages of approximately twenty-five repeats[1]. Figures 3.8 and 3.9 show example EEG and [ΔHbR] signals that Combined DCM was fitted to.

To improve the Combined DCM fits, it was necessary to pre-process the [ΔHbR] signals. Specifically, because the [ΔHbR] responses often did not return to baseline before the next train of stimuli was delivered (the average inter-stimulus interval of the dataset is only 12s), it was necessary to extend the responses being fitted to in order to prevent the fits to the main responses being affected detrimentally (resulting from the vasodilation differential equation being a damped

[1]The disadvantage of only fitting Combined DCM to the average responses for each brain hemisphere is reduced statistical power for any inferences made about the Combined DCM parameters because of low sample size (approximately only ten datapoints per anaesthetic). This will prove problematic for the work in Chapter 4.

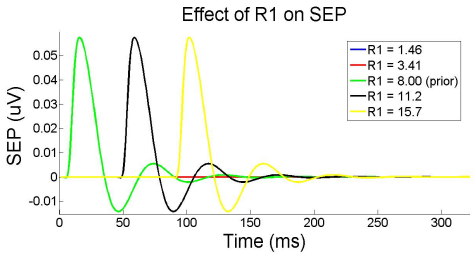
Effect of R1 on SEP
R1 = 1.46
R1 = 3.41
R1 = 8.00 (prior)
R1 = 11.2
R1 = 15.7
SEP (uV)
Time (ms)

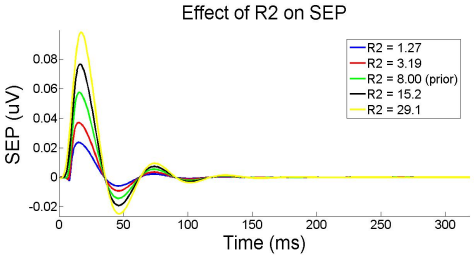
Effect of R2 on SEP
R2 = 1.27
R2 = 3.19
R2 = 8.00 (prior)
R2 = 15.2
R2 = 29.1
SEP (uV)
Time (ms)

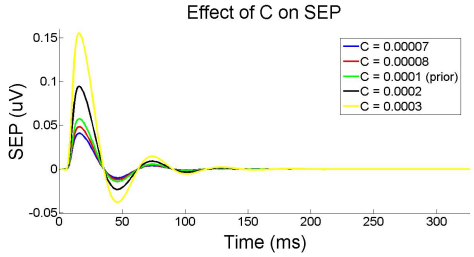
Effect of C on SEP
C = 0.00007
C = 0.00008
C = 0.0001 (prior)
C = 0.0002
C = 0.0003
SEP (uV)
Time (ms)

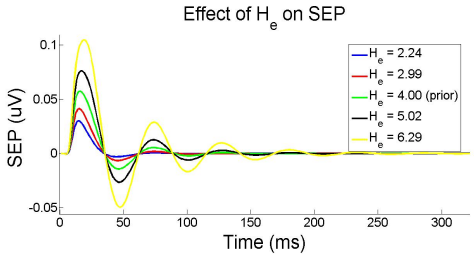
Effect of H_e on SEP
H_e = 2.24
H_e = 2.99
H_e = 4.00 (prior)
H_e = 5.02
H_e = 6.29
SEP (uV)
Time (ms)

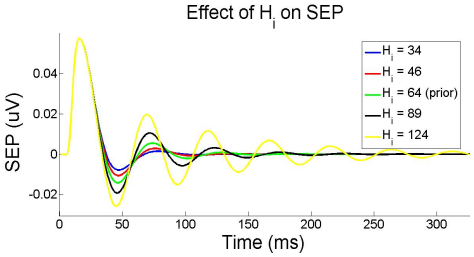
Effect of H_i on SEP
H_i = 34
H_i = 46
H_i = 64 (prior)
H_i = 89
H_i = 124
SEP (uV)
Time (ms)

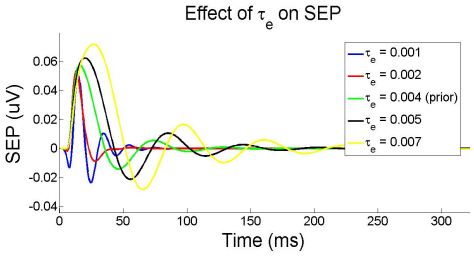
Effect of τ_e on SEP
τ_e = 0.001
τ_e = 0.002
τ_e = 0.004 (prior)
τ_e = 0.005
τ_e = 0.007
SEP (uV)
Time (ms)

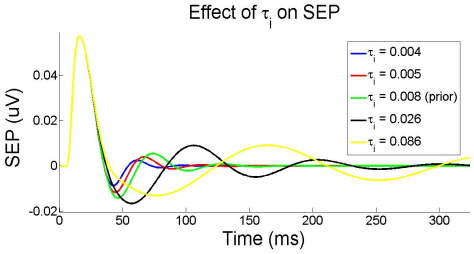
Effect of τ_i on SEP
τ_i = 0.004
τ_i = 0.005
τ_i = 0.008 (prior)
τ_i = 0.026
τ_i = 0.086
SEP (uV)
Time (ms)

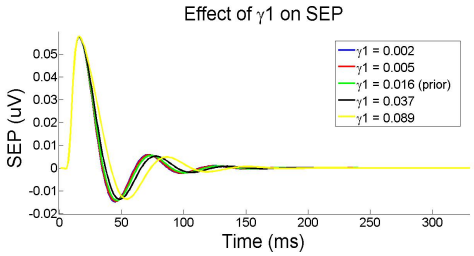
Effect of γ1 on SEP
γ1 = 0.002
γ1 = 0.005
γ1 = 0.016 (prior)
γ1 = 0.037
γ1 = 0.089
SEP (uV)
Time (ms)

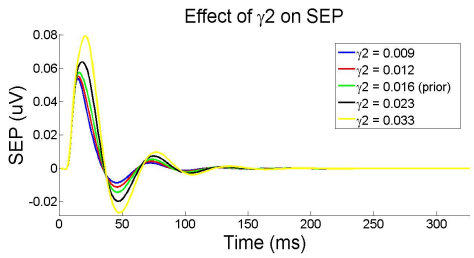
Effect of γ2 on SEP
γ2 = 0.009
γ2 = 0.012
γ2 = 0.016 (prior)
γ2 = 0.023
γ2 = 0.033
SEP (uV)
Time (ms)

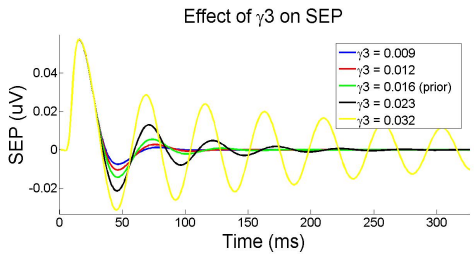
Effect of γ3 on SEP
γ3 = 0.009
γ3 = 0.012
γ3 = 0.016 (prior)
γ3 = 0.023
γ3 = 0.032
SEP (uV)
Time (ms)

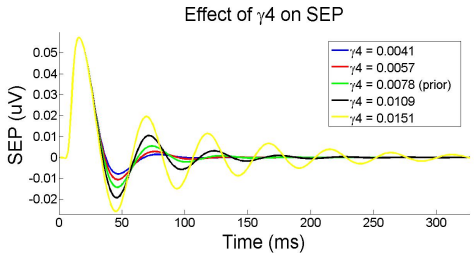
Effect of γ4 on SEP
γ4 = 0.0041
γ4 = 0.0057
γ4 = 0.0078 (prior)
γ4 = 0.0109
γ4 = 0.0151
SEP (uV)
Time (ms)

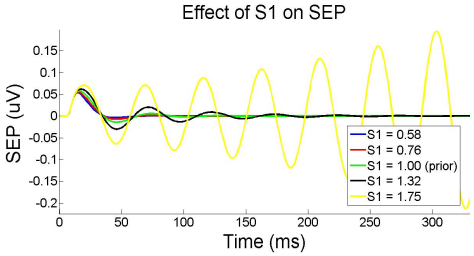
Effect of S1 on SEP
S1 = 0.58
S1 = 0.76
S1 = 1.00 (prior)
S1 = 1.32
S1 = 1.75
SEP (uV)
Time (ms)

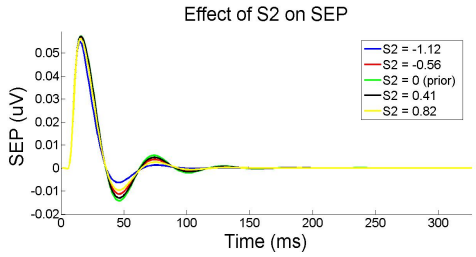

Figure 3.6: SEP responses for each DCM-EEG parameter varied between its minimum and maximum posterior values from the DCM-EEG fits to the real per-rat brain hemisphere EEG data.

oscillator and therefore having infinite power). The pre-processing steps were:

1. Fit a sinusoid to the final second of the true signal, and then extend the timeseries so that it is 29 seconds long using this sinusoid.
2. Exponentially decay the extended signal so that it returns to baseline by the end.

As the longest signal in the whole dataset was 29 seconds, these pre-processing steps meant that no signal potentially carrying information about anaesthetic condition was ignored.

In an attempt to speed up the per-rat brain hemisphere fittings, the prior parameter moments listed in Table 3.1 were used to first fit Combined DCM to the mean responses to each anaesthetic, which are considerably less noisy than the per-rat brain hemisphere responses. The posterior parameter mean values from these per-anaesthetic Combined DCM fittings were then used as the prior parameter means for the per-rat brain hemisphere fittings.

To balance the accuracies of the Combined DCM fittings to the EEG and [ΔHbR] signals, Combined DCM was fitted to the [ΔHbR] signal and then the resulting EEG-related posteriors were used as priors to a DCM-EEG fitting to the EEG signal only at each E-M iteration. The DCM-EEG posteriors then served as priors for the EEG-related parameters in Combined DCM in the next iteration, but with considerably tighter prior covariances for the DCM-EEG parameters

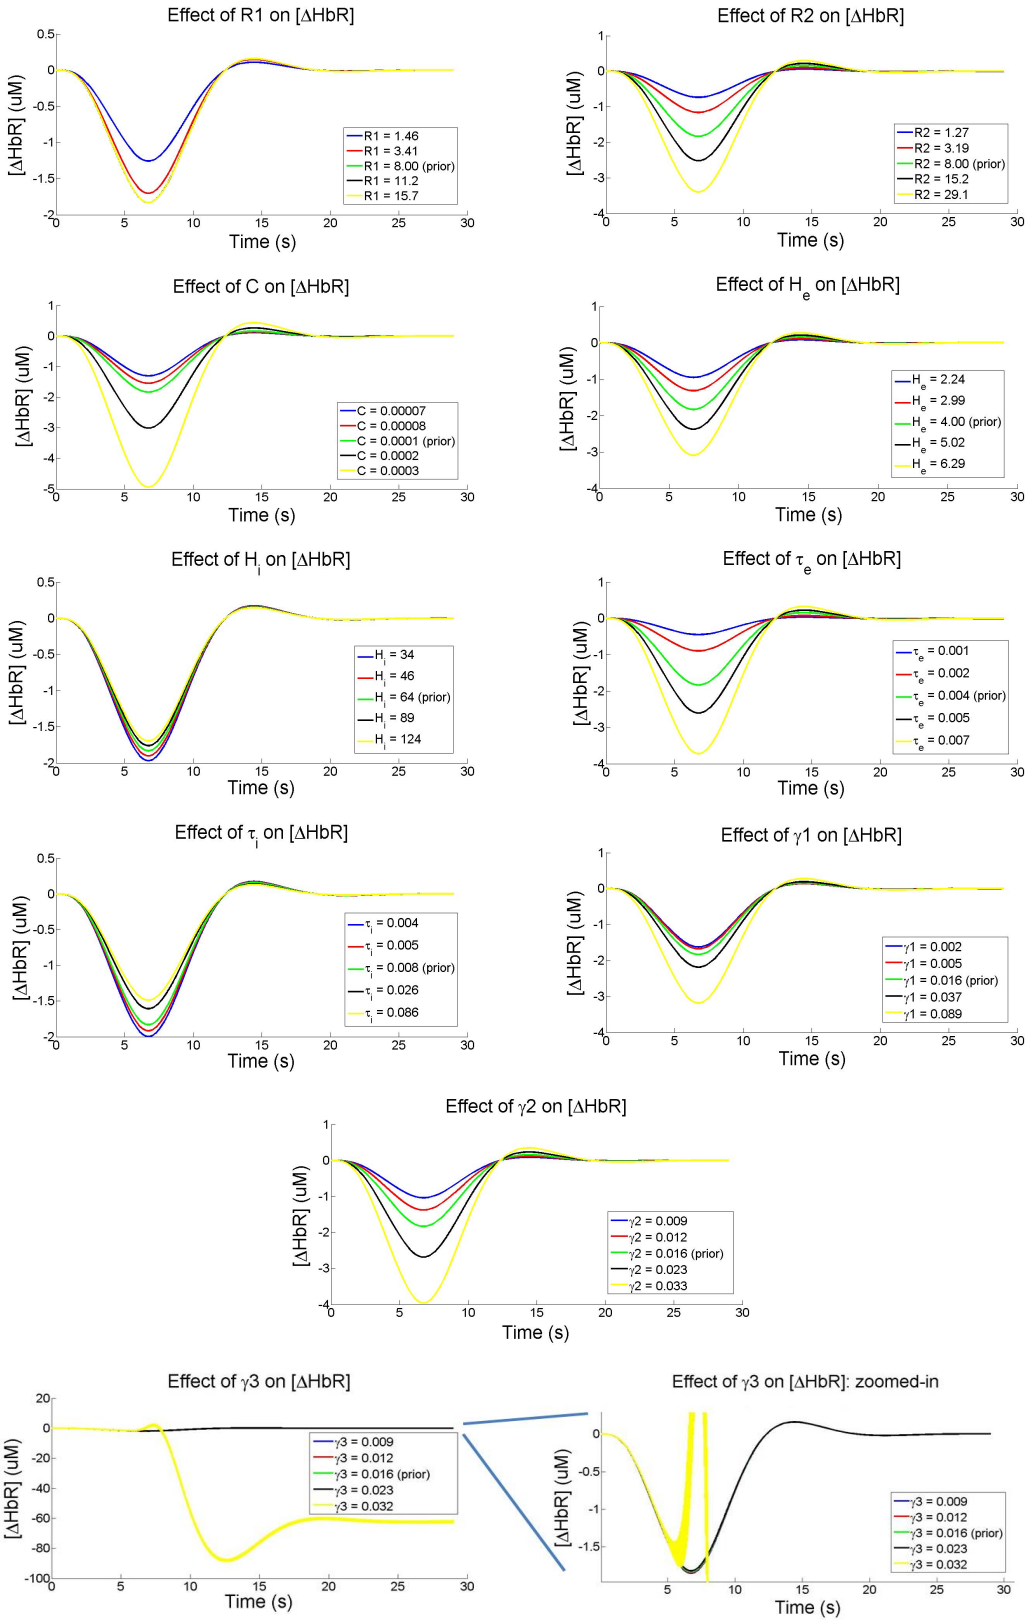
Effect of R1 on [ΔHbR]
[ΔHbR] (uM)
Time (s)
R1 = 1.46
R1 = 3.41
R1 = 8.00 (prior)
R1 = 11.2
R1 = 15.7
Effect of R2 on [ΔHbR]
[ΔHbR] (uM)
Time (s)
R2 = 1.27
R2 = 3.19
R2 = 8.00 (prior)
R2 = 15.2
R2 = 29.1
Effect of C on [ΔHbR]
[ΔHbR] (uM)
Time (s)
C = 0.00007
C = 0.00008
C = 0.0001 (prior)
C = 0.0002
C = 0.0003
Effect of H_e on [ΔHbR]
[ΔHbR] (uM)
Time (s)
H_e = 2.24
H_e = 2.99
H_e = 4.00 (prior)
H_e = 5.02
H_e = 6.29
Effect of H_i on [ΔHbR]
[ΔHbR] (uM)
Time (s)
H_i = 34
H_i = 46
H_i = 64 (prior)
H_i = 89
H_i = 124
Effect of τ_e on [ΔHbR]
[ΔHbR] (uM)
Time (s)
τ_e = 0.001
τ_e = 0.002
τ_e = 0.004 (prior)
τ_e = 0.005
τ_e = 0.007
Effect of τ_i on [ΔHbR]
[ΔHbR] (uM)
Time (s)
τ_i = 0.004
τ_i = 0.005
τ_i = 0.008 (prior)
τ_i = 0.026
τ_i = 0.086
Effect of γ1 on [ΔHbR]
[ΔHbR] (uM)
Time (s)
γ1 = 0.002
γ1 = 0.005
γ1 = 0.016 (prior)
γ1 = 0.037
γ1 = 0.089
Effect of γ2 on [ΔHbR]
[ΔHbR] (uM)
Time (s)
γ2 = 0.009
γ2 = 0.012
γ2 = 0.016 (prior)
γ2 = 0.023
γ2 = 0.033
Effect of γ3 on [ΔHbR]
[ΔHbR] (uM)
Time (s)
γ3 = 0.009
γ3 = 0.012
γ3 = 0.016 (prior)
γ3 = 0.023
γ3 = 0.032
Effect of γ3 on [ΔHbR]: zoomed-in
[ΔHbR] (uM)
Time (s)
γ3 = 0.009
γ3 = 0.012
γ3 = 0.016 (prior)
γ3 = 0.023
γ3 = 0.032

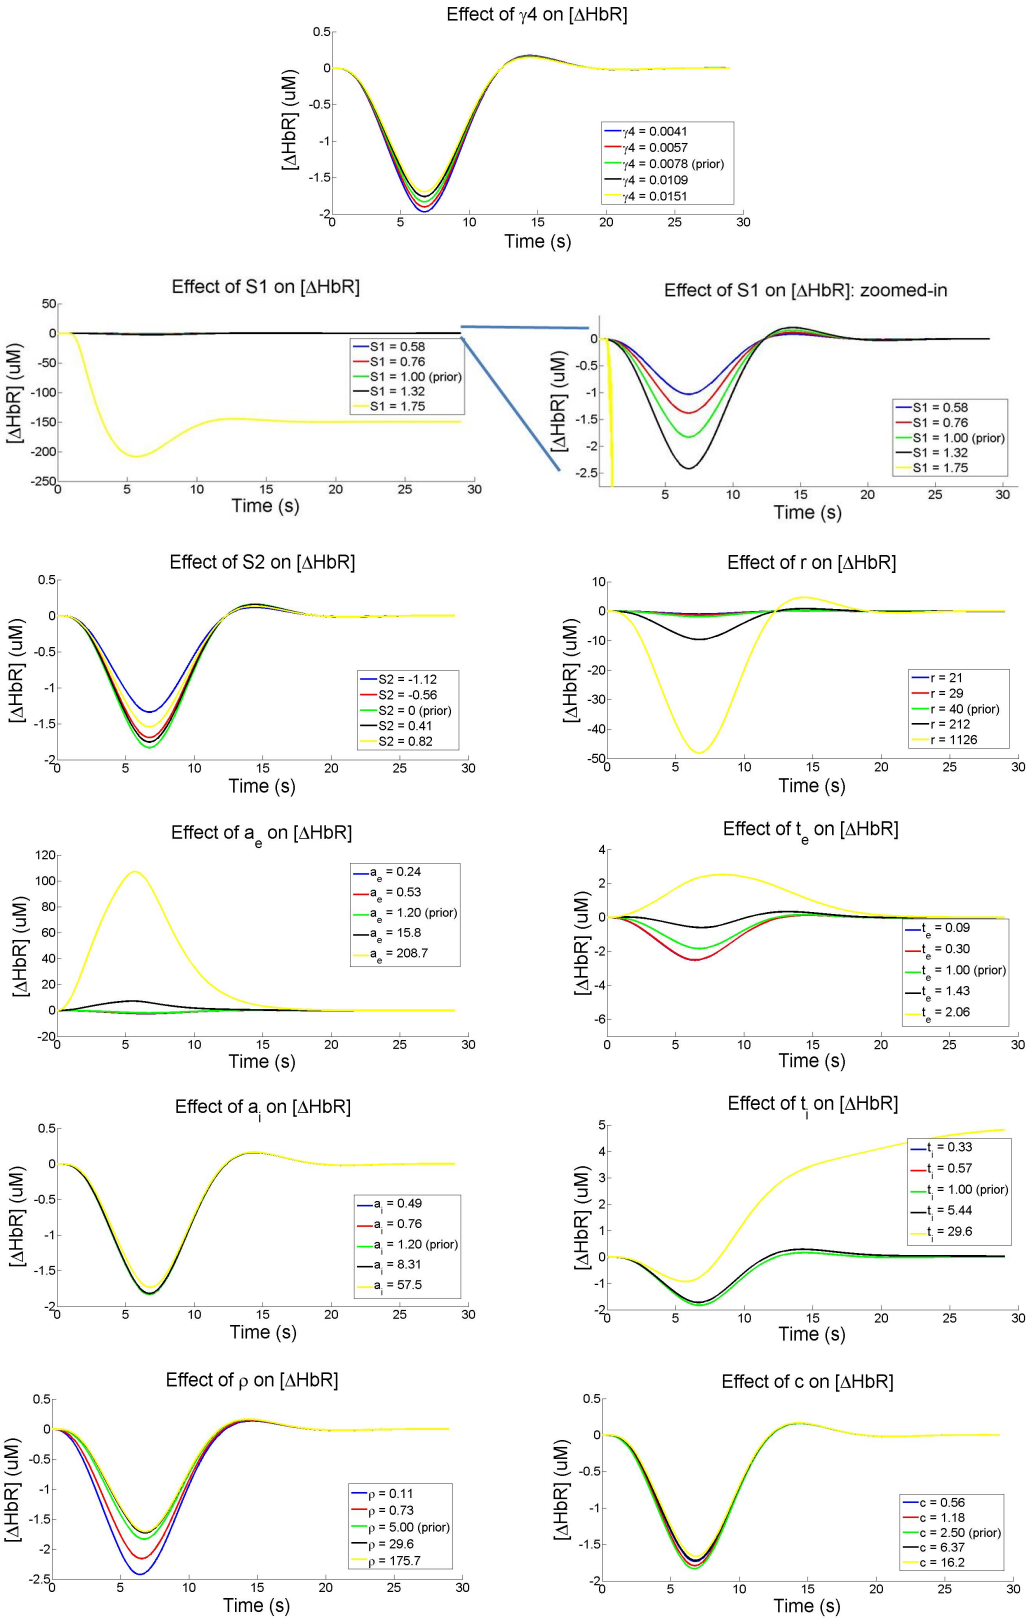
Effect of γ4 on [ΔHbR]
[ΔHbR] (uM)
Time (s)
γ4 = 0.0041
γ4 = 0.0057
γ4 = 0.0078 (prior)
γ4 = 0.0109
γ4 = 0.0151
Effect of S1 on [ΔHbR]
Effect of S1 on [ΔHbR]: zoomed-in
S1 = 0.58
S1 = 0.76
S1 = 1.00 (prior)
S1 = 1.32
S1 = 1.75
Effect of S2 on [ΔHbR]
S2 = -1.12
S2 = -0.56
S2 = 0 (prior)
S2 = 0.41
S2 = 0.82
Effect of r on [ΔHbR]
r = 21
r = 29
r = 40 (prior)
r = 212
r = 1126
Effect of a_e on [ΔHbR]
a_e = 0.24
a_e = 0.53
a_e = 1.20 (prior)
a_e = 15.8
a_e = 208.7
Effect of t_e on [ΔHbR]
t_e = 0.09
t_e = 0.30
t_e = 1.00 (prior)
t_e = 1.43
t_e = 2.06
Effect of a_i on [ΔHbR]
a_i = 0.49
a_i = 0.76
a_i = 1.20 (prior)
a_i = 8.31
a_i = 57.5
Effect of t_i on [ΔHbR]
t_i = 0.33
t_i = 0.57
t_i = 1.00 (prior)
t_i = 5.44
t_i = 29.6
Effect of ρ on [ΔHbR]
ρ = 0.11
ρ = 0.73
ρ = 5.00 (prior)
ρ = 29.6
ρ = 175.7
Effect of c on [ΔHbR]
c = 0.56
c = 1.18
c = 2.50 (prior)
c = 6.37
c = 16.2

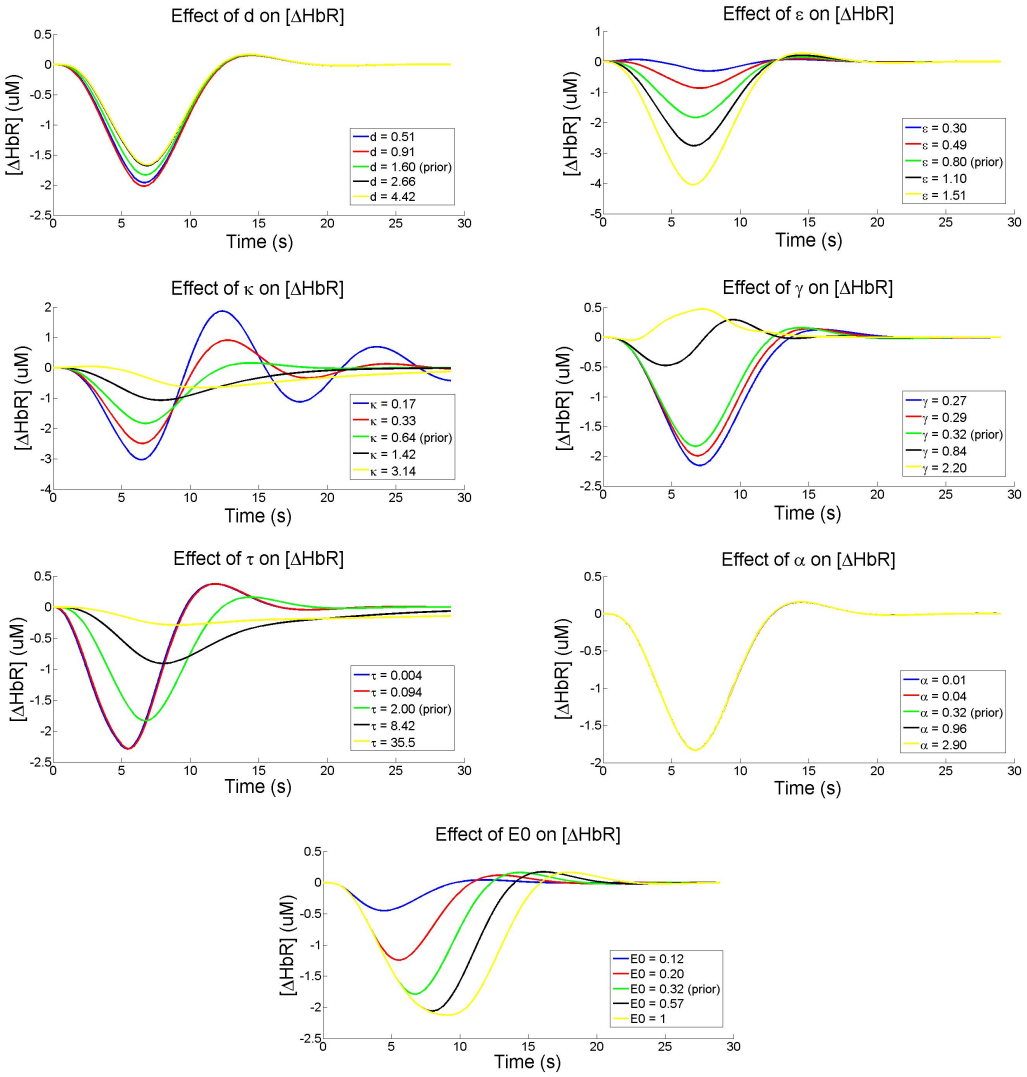

Figure 3.7: [ΔHbR] responses for each Combined DCM parameter varied between its minimum and maximum posterior values from the Combined DCM fits to the real per-rat brain hemisphere [ΔHbR] data.

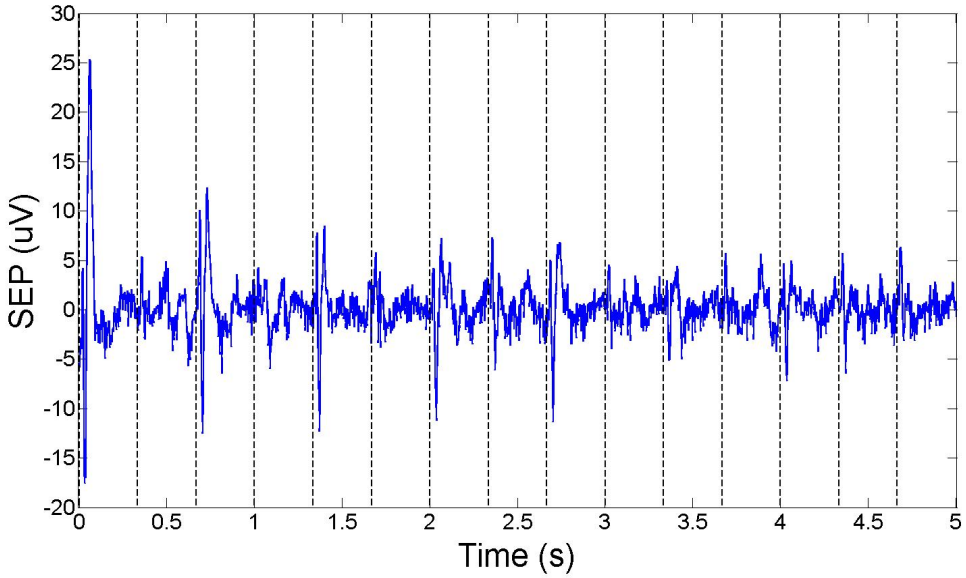

Figure 3.8: Example (Alpha-Chloralose) EEG timeseries that Combined DCM was fitted to. This timeseries shows an EEG response to a 5s, 3Hz train of electrical stimuli (15 responses), occuring at the dashed lines. No pre-processing was done to these data.

($\frac{1}{100}^{th}$ their covariances listed in Table 3.1) to prevent the [ΔHbR] fit pushing these DCM-EEG parameters too far from their posterior values from the DCM-EEG fitting. This iteration continued until both the free energy of the Combined DCM fitting to the [ΔHbR] signal and the free energy of DCM-EEG to the EEG signal converged.

3.2.2 Simulating [ΔHbR] data

The posteriors resulting from Combined DCM fitting to the real data were used to simulate [ΔHbR] timeseries similar to the real [ΔHbR] timeseries except for having "unanaesthetised" hemodynamics. This was achieved using Cholesky decomposition to sample from the m[1] multivariate Gaussians, $\mathbf{Y}_{mx27}$, that approximate the posterior parameter spaces, $\mathbf{X}_{mx27}$, for each anaesthetic n[2] times.

[1] m is equal to the number of mean responses to 5s stimulation trains that Combined DCM is fitted to under each anaesthetic, which is 55 for Alpha-Chloralose, 50 for Pentobarbital, 43 for Ketamine, 55 for Fentanyl, 50 for Isoflurane and 60 for Propofol.

[2] n is equal to the number of stimulation trains (of any of the seven stimulation durations) that occurred during a single run, ranging from 22 to 35.

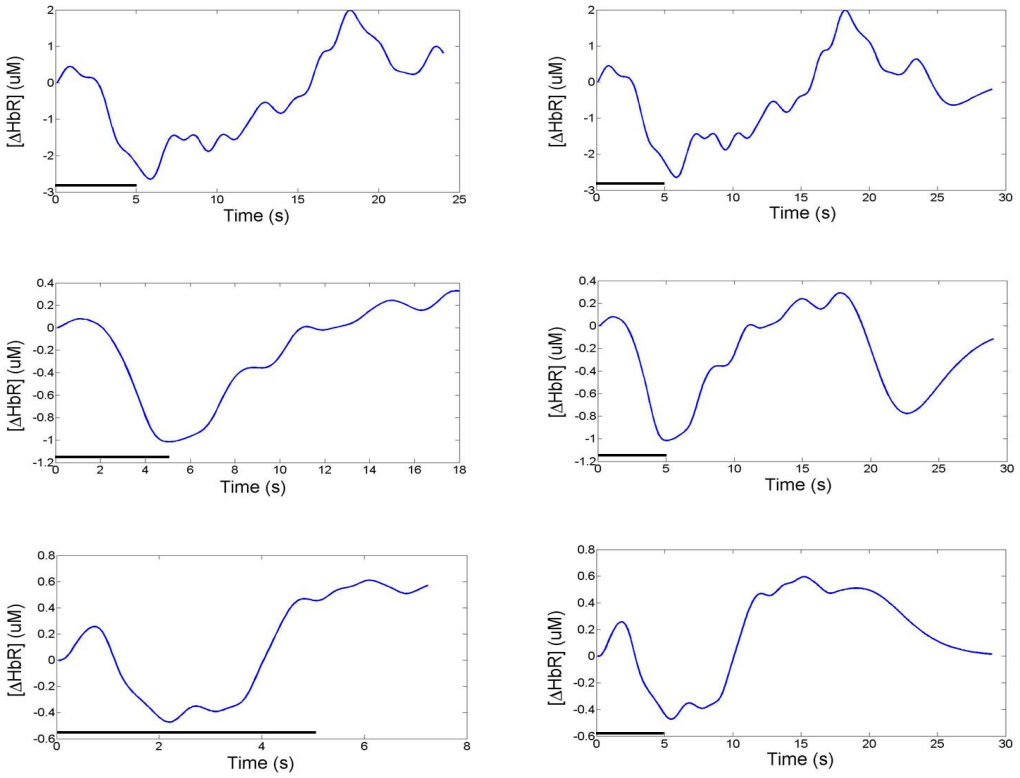

(a) Before pre-processing

(b) After pre-processing

Figure 3.9: Three example (Alpha-Chloralose) [ΔHbR] timeseries (a) before pre-processing, and (b) after pre-processing (which Combined DCMs were fitted to).

$$\mathbf{Y} = \mathbf{Ep} + \mathbf{X}\mathbf{U}_{Cp}\,\mathrm{diag}(\mathbf{Cp}^{1/2}), \tag{3.12}$$

where $\mathbf{Ep}$ = 27 posterior means, with $\mathbf{Ep}_{EEG}$ taking the anaesthetised per-rat brain hemisphere values but $\mathbf{Ep}_{fMRI}$ taking canonical [ΔHbR] values, $\mathbf{Cp}$ = posterior covariances between the 27 parameters, with $\mathbf{Cp}_{EEG}$ and $\mathbf{Cp}_{EEG-fMRI}$ taking the anaesthetised per-rat brain hemisphere values but $\mathbf{Cp}_{fMRI}$ taking canonical [ΔHbR] values, and $\mathbf{U}_{Cp}$ = Cholesky decomposition of $\mathbf{Cp}$ ($\mathbf{Cp} = \mathbf{U}_{Cp}\mathbf{U}_{Cp}^{*}$, where * = conjugate transpose)[1].

To complete the [ΔHbR] data simulation, real [ΔHbR] background signals (the resting-state [ΔHbR] signals occuring before stimuli were administered) were added to the simulated signals to provide realistic noise.

3.2.3 Embedding and Inference

Finally, GLM identical to that carried out on the real [ΔHbR] signals in Chapter 2 was then carried out on these simulated data and the significances of the differences in their Z-values and their different reliabilities were compared to those calculated for the real data.

3.3 Results

Tables 3.2 and 3.3 report the significances and reliabilities with which the GLM framework employed by SPM8 can detect neural activations in [ΔHbR] signals simulated from anaesthetic-specific neuronal activities but "unanaesthetised" hemodynamics. Sample GLM fits for each of the anaesthetics are shown in Fig.3.10 (these samples are the same as those shown in Fig.2.7 - 2.12 in Chapter 2 for the real data, making these two sets of figures directly comparable).

[1]Notice that, as the metabolism parameters in Sotero et al.'s model control aspects of the [ΔHbR] shape rather than just its amplitude (as shown in Fig.3.7), these parameters are considered "hemodynamic" parameters and therefore grouped with the fMRI parameters, i.e. the fMRI parameters consist of both the Sotero et al. and mDCM-fMRI parameters.

Anaesthetic	**Z value using canonical [ΔHbR]**			
		+ temporal derivative	**+ dispersion derivative**	**+ both derivatives**
Alpha-Chloralose	7.0597*** †††	7.0986*** †††	7.1532*** †††	7.1791*** †††
Pentobarbital	6.4200*** †††	6.4059*** †††	6.3562*** †††	6.3555*** †††
Ketamine	2.4174** †	2.4209** †	2.4160** †	2.4208** †
Fentanyl	4.6857*** †	4.7069*** †	4.7006*** †	4.7155*** †
Isoflurane	2.9658**	3.0113**	2.8453**	2.8576**
Propofol	6.1730*** †††	6.5149*** †††	6.1251*** †††	6.1165*** †††

Table 3.2: Multi-subject neural activation Z-values for [ΔHbR] data simulated under each anaesthetic using GLM with the canonical hemodynamic response function and different combinations of its temporal and dispersion derivatives. * indicates that a Z-value equivalent to $p < 0.05$, ** equivalent to $p < 0.01$ and *** equivalent to $p < 0.001$. † indicates that a Z-value for the simulated data is significantly different from the corresponding Z-value for the real data in Chapter 2 with $p < 0.05$, †† with $p < 0.01$ and ††† with $p < 0.001$

Anaesthetic	**Percentage of activations detected as significant value using canonical [ΔHbR]**			
		+ temporal derivative	**+ dispersion derivative**	**+ both derivatives**
Alpha-Chloralose	100%	100%	100%	100%
Pentobarbital	100%	100%	100%	100%
Ketamine	100%	100%	100%	100%
Fentanyl	100%	100%	100%	100%
Isoflurane	70%	70%	70%	70%
Propofol	100%	100%	100%	100%

Table 3.3: Percentage of rat brain hemispheres detected to have significant (equivalent to $p < 0.05$) neural activation from [ΔHbR] signals simulated under each anaesthetic using GLM with the canonical hemodynamic response function different combinations of its temporal and dispersion derivatives.

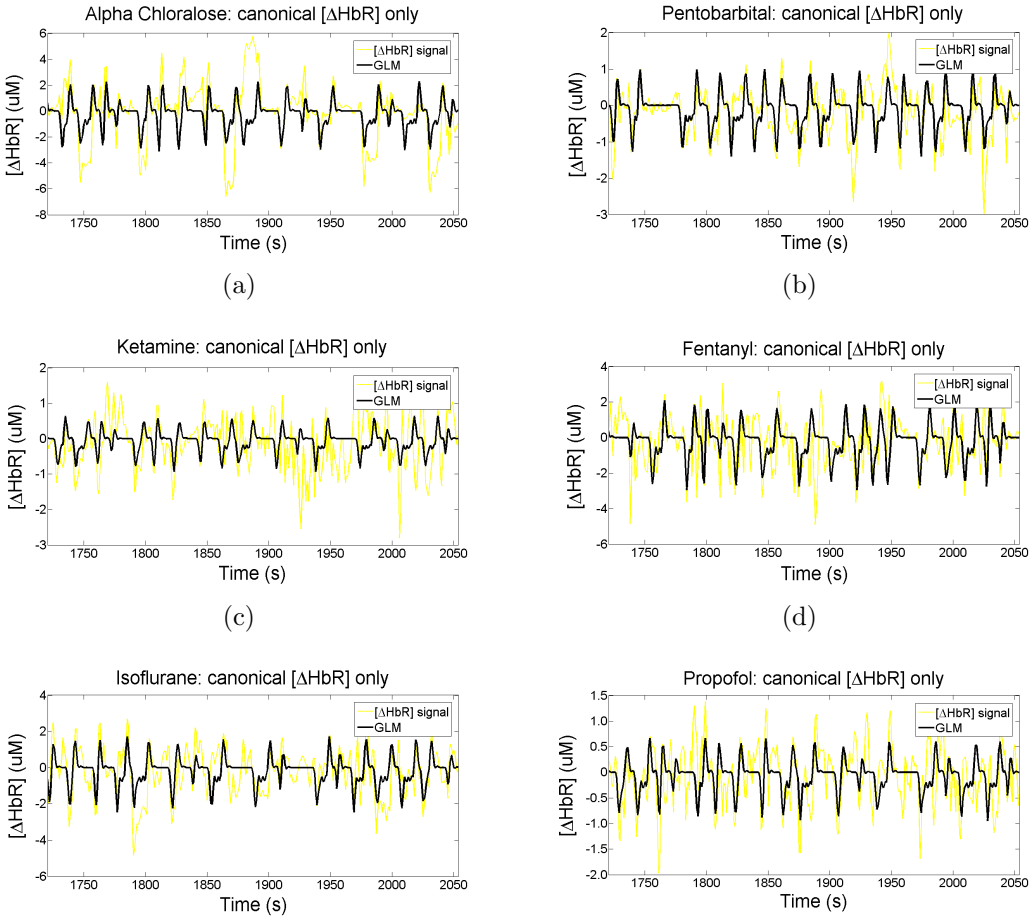

Figure 3.10: GLM fits for example simulated [Δ HbR] signal under each anaesthetic when only canonical [ΔHbR] is included as a basis function in the GLM (adding the temporal and/or dispersion derivatives makes little difference to these fits). The Z-values for these example are (a) 34.4146***, (b) 37.1238***, (c) 25.6622***, (d) 27.3022***, (e) 36.7735*** and (f) 29.4514***.

As the real and simulated [ΔHbR] signals were driven by similar neuronal activity[1], the hypothesis that aspects of neurovascular coupling relevant to BOLD GLM are unaltered by anaesthetic modulations can only be true if the Z-values in Table 3.2 are similar to those calculated in Chapter 2 for the real data (listed in Table 2.2). Two-sample paired t-tests reveal that there are significant differences in the individual Z-values for the real and simulated signals under all of the anaesthetics except for Isoflurane, as indicated in Table 3.2. However, the finding of greater consequence is that that there are differences in the significance **levels** of the group-level Z statistics under Pentobarbital, Ketamine, Isoflurane and Propofol: the effects of Ketamine and Isoflurane on hemodynamics lead to overestimations of the underlying neuronal activity (it is detected as being extremely significant when it fact it is only very significant)[2], whereas the effects of Pentobarbital and Propofol on hemodynamics lead to underestimations of the underlying neuronal activity (no significant neuronal activity is detected under Pentobarbital, when it is in fact extremely significant, and only very significant neuronal activity is detected under Propofol when it too is actually extremely significant). The only anaesthetics under which neuronal activity was detected from the real hemodynamics responses at correct significance levels are Alpha-Chloralose and Fentanyl. The discrepancies in the group-level Z values are much smaller under Fentanyl than under Alpha-Chloralose (-0.4399 compared to 2.1171 when only canonical [ΔHbR] is included as a basis function).

[1]Combined DCM is a model, i.e. a simplified formalisation of reality that balances accuracy and practicality, and this means that the neuronal activity used to generate the simulated hemodynamics responses will not be identical to the real neuronal activity that generated the real hemodynamics responses. The tight covariances of the DCM-EEG parameters of the Combined DCM fitting will have ensured that the neuronal activity driving both the real and simulated hemodynamics are at least similar, however. The limited accuracy of Combined DCM is explored in the Conclusions chapter.

[2]Three runs under Isoflurane were detected as significant despite not actually being significant at all (compare Tables 2.3 and 3.3).

3.4 Conclusions

Chapters 2 and 3 have shown that BOLD studies investigating the effect of anaesthesia on neuronal activity (i.e. preclinical phfMRI studies in which anaesthetic is the interest neuromodulation, or parametric neurovascular coupling experiments) and analysed using standard GLM assumptions are confounded when they are conducted under Pentobarbital, Ketamine, Isoflurane or Propofol anaesthesia. If Ketamine or Isoflurane is used, this work suggests that the resulting Z values will overestimate the significance of the underlying neuronal activity, whereas if Pentobarbital or Propofol is used, this work suggests that the resulting Z values will underestimate the significance of the underlying neuronal activity. The repercussions of these confounding effects for the conclusions made about group-level neuronal activity levels from standard BOLD experiments conducted under Ketamine, Isoflurane and Propofol are likely to be a relatively modest, but they may be dramatic for experiments conducted under Pentobarbital as the effects of this anaesthetic on neurovascular coupling are highly detrimental to GLM significance (extremely significant group-level neuronal activity goes completely undetected).

This work has also shown that BOLD studies investigating the effects of Alpha-Chloralose or Fentanyl on neuronal activity are not confounded when standard GLM assumptions are adopted. As the discrepancies in the group-level Z-values are much smaller under Fentanyl than under Alpha-Chloralose, the main finding therefore is that Fentanyl is the best anaesthetic for these types of experiments.

More generally, Chapters 2 and 3 have shown that any BOLD-fMRI experiment may be confounded by interest and/or nuisance neuromodulations, whether they are pharmacological agents other than anaesthetics, or brain pathologies. This chapter has demonstrated that the generative embedding scheme outlined in Fig.1.4 provides a way for these BOLD studies to be identified.

The question is now what can be done to prevent BOLD-fMRI studies (where neuronal activity is not measured) from being confounded? A number of methods described in the literature have been developed to reduce the inaccuracies of BOLD GLM analyses. These methods fall into two groups: 1) those that enable the nonlinearities in the neuronal activity and/or hemodynamics to be captured, and 2) those that tailor the impulse response function, reflecting the two ways in which neuromodulations can alter hemodynamics to the detriment of BOLD GLM analysis validity:

1. By altering the hemodynamic response function to a single, short stimulus (the hemodynamic impulse response).
2. By altering the linearity of the hemodynamic response that the GLM models[1].

This chapter concludes with a description of the value that generative embedding could bring to some of these methods to make them more practical to use.

3.4.1 Capturing nonlinearities in GLM

The most obvious way to capture nonlinearities is to fit a biophysical model such as the Balloon Model or mDCM-fMRI to the hemodynamics data. The three major problems with this approach are 1) it is very computationally intensive to fit these sorts of models to long timeseries of data involving multiple responses, 2) they may suffer convergence problems, and 3) there is no clear and obvious way to

[1]It is possible to have a separate explanatory variable for each stimulation paradigm, i.e. to use first-order response functions to create separate explanatory variables for each stimulation train duration, however this approach does not completely mitigate the problem of nonlinearity because it is biased towards the more linear responses that are better captured by the first-order responses. This approach also brings with it the issue of interpretability when the focus of the study is not to understand differences in neuronal activity arising from different stimulation paradigms as it results in a separate significance level for each stimulation paradigm.

calculate significance levels for the parameters of interest (analagous to the $\boldsymbol{\beta}$ in BOLD GLMs) that researchers are familiar with[1].

An alternative method for capturing nonlinearities that suffers from these problems to a much lesser extent is the Volterra kernels formulation developed by Friston et al., 1998 [62]. A Taylor expansion of Volterra kernels can completely characterise the dynamics of a time-invariant system:

$$\begin{aligned} y(t) = h^0 &+ \int_{-\infty}^{+\infty} h^1(\tau_1).u(t-\tau_1)d\tau_1 \\ &+ \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} h^2(\tau_1,\tau_2).u(t-\tau_1).u(t-\tau_2)d\tau_1 d\tau_2 + ... \\ &+ \int_{-\infty}^{+\infty} ... \int_{-\infty}^{+\infty} h^n(\tau_1,...,\tau_n).u(t-\tau_1)....u(t-\tau_n)d\tau_1...d\tau_n + ..., \end{aligned} \tag{3.13}$$

where $y(t)$ is the output timeseries of the system being characterised (here hemodynamic signal), $u(t)$ is the stimulation input at time t, and $\mathrm{h}^n(\tau_1,...,\tau_n)$ is the n^{th}-order Volterra kernel.

Friston et al. reformulated Equ.3.13 into a GLM so that it could be estimated using least squares and therefore analysed using the traditional ANOVA statistics familiar to current BOLD GLM users. This was achieved by capturing all of the nonlinearities in the explanatory variables, $\boldsymbol{x}(t)$, allowing their coefficients, $\mathbf{g}$, to remain linear.

$$\begin{aligned} y(t) = g^0 &+ \int_{i=1}^{P} g_i^1 x_i(t) \\ &+ \int_{i=1}^{P}\int_{j=1}^{P} g_{ij}^1 x_i(t)\cdot x_j(t) + ... \\ &+ \int_{i=1}^{P} ... \int_{n=1}^{P} g_{i..n}^n x_i(t)\cdot ...\cdot x_n(t) + ..., \end{aligned} \tag{3.14}$$

[1]Biophysical model fitting could be considerably sped up if the average response to a stimulus type were fitted to rather than the full timeseries, but this complicates the interpretation of statistical significance further because the standard error of the average fit would be a very large underestimation of the standard error of the model fit to the raw timeseries.

where $\boldsymbol{x}(t) = \int \boldsymbol{b}(\tau_1) \cdot u(t - \tau_1) d\tau_1$, i.e. the first-order term $x_i(t)$ is the convolution of the basis functions $\boldsymbol{b}_i$ (three gamma functions in Friston et al., 1998) with the input u, the second-order term $x_j(t)$ is the convolution of the first temporal derivatives of $\boldsymbol{b}_i$ with u, etc.

Fitting Volterra kernels is computationally cheaper than fitting biophysical models when the Taylor series is limited to the first few Volterra kernels, and Friston et al. found that when the Taylor expansion was limited to second order, i.e. the first two lines of Equ.3.14, the most important nonlinear behaviours (the interaction of the hemodynamic response to one stimulus given the previous stimuli) were still captured.

Unfortunately, although Volterra kernels can be embedded into the same statistics as traditional BOLD GLMs, it is not possible to estimate the signal attributable to the parameter of interest only because of collinearity between the first- and second-order kernels.

Wagner et al., 2005, have developed a method for fitting nonlinear models that is also computationally faster than fitting biophysical models but with which the parameter of interest can be reliably estimated [192]. In this method, low-dimensional nonlinearities are introduced by allowing the response amplitude, onset time and time-to-peak shape features of the GLM regressors to vary as empirically derived biexponential functions of stimulation history. However, limiting the dimensionality of the problem in this way comes at the expense of not being able to capture nonlinearities in aspects of response shape other than amplitude, onset time and time-to-peak.

3.4.2 Tailoring HRFs in GLM

The literature regarding ways to tailor the HRF in GLMs is more developed than the literature regarding ways to capture nonlinearities in GLMs. The most common method for tailoring the HRF is to use the first few principal components

of the data as basis functions, because they capture the hemodynamic behaviour most efficiently [79]. In Hossein-Zadeh et al., they found that the first three principal components captured 99% of the variance in their BOLD dataset, with the first principal component closely approximating a Gamma function. However, they found clear differences between the second principal component and the temporal derivative of the Gamma function, and between the third principal component and the dispersion derivative of the Gamma function, indicating the limited benefit of including the canonical HRF derivatives in GLM analyses [79].

The disadvantage of using principal components as basis functions is that they have very high flexibility, meaning that noise may be fitted to as well as response. One way to lessen this problem is to collect an extra dataset and calculate the basis functions from these data rather than the experimental data [3], however the obvious disadvantage of this method is that it increases the financial expense of a BOLD experiment as more scanning time is required. An alternative and more efficacious method is to constrain the basis functions algorithmically. Woolrich et al., 2004, developed a method to constrain basis functions so that only biologically plausible hemodynamics could be modelled by the GLM [198]. This is achieved by first identifying an allowable subspace of basis functions (approximated by a multivariate normal) from the principal components of biologically plausible hemodynamic responses, and then using Variational Bayes to impose soft constraints on the proposed basis functions during fitting. Unfortunately, this approach is relatively computationally expensive.

3.4.3 Correcting for neuromodulations using generative embedding

One immediately implementable way to account for modulations of hemodynamic linearity and impulse response shape in analyses of BOLD data without complicating the simple (ANOVA) reporting statistics would be to apply both

Wagner et al.'s nonlinearity correction and Woolrich et al.'s impuse response shape correction to the GLM regressors. As already mentioned, these algorithms have weaknesses that prevent them from being wholly effective or practical. Generative embedding provides a way to overcome these weaknesses by identifying which subset(s) of the generative model parameters should:

1. capture the nonlinear behaviours of the hemodynamics more fully than response amplitude, onset time and time-to-peak shape features, allowing (biexponential) functions to be fitted to these parameters in Wagner et al.'s nonlinearity correction instead.
2. sufficiently explain the variation in the allowable subspace of basis functions in Woolrich et al.'s impulse response shape correction, allowing this subspace to be constrained even more and, in doing so, speeding up the GLM fitting.

The aim of Chapter 4 will be to use generative modelling and generative embedding to identify what these generative model parameters might be in the context of anaesthetic neuromodulations.

Chapter 4

Insights into the effects of anaesthetic neuromodulation on neurovascular coupling

In Chapter 3 it was found that the extent to which anaesthetics confound GLM by altering neurovascular coupling varies across anaesthetics. This chapter describes efforts to shed light on how anaesthetics cause these GLM significance discrepancies.

As already mentioned, there are two ways in which neuromodulations can alter hemodynamics to the detriment of BOLD GLM significances:

1. By altering the shape of the hemodynamic response function to a single, short stimulus (the hemodynamic impulse response),
2. By altering the linearity of the hemodynamic response.

The aim of the first half of this chapter is to make qualitative conclusions about the contributions of these two phenomena to the Z-value discrepancies caused by the six anaesthetics. This is achieved using generative modelling.

In the second half of this chapter, generative embedding is used to identify

which specific biological actions of the anaesthetics explain these discrepancies[1].

4.1 Qualitative observations

Even without implementing a sophisticated method to separate the contributions of anaesthetic modulations of hemodynamic impuse response and linearity to Z-value discrepancies, it is possible to make high-level statements about which of these phenomena contribute strongly to the discrepancies under each anaesthetic by simply comparing them (see Fig.4.1 - 4.6). This is possible because the impulse response shape and amplitude remain constant for all stimulation train durations, meaning that any patterns of differences between 1) the GLM "responses" fitted to the simulated (unanaesthetised) data, and 2) the real responses, across the stimulation train durations must be attributable to hemodynamic nonlinearities (assuming that Combined DCM is largely accurate). Consistent differences across the stimulation train durations must instead be explained by mismatch of canonical [ΔHbR] shape to impulse response shape.

The most obvious observation to make from Fig.4.1 - 4.6 is that there are large discrepancies in the amplitudes of the average impulse responses under Alpha-Chloralose, Pentobarbital, Isoflurane and Propofol compared to the average GLM responses[2]. These discrepancies result from the shapes of these impulse responses being very different from the shape of canonical [ΔHbR] - generally much slower.

In contradistinction, canonical [ΔHbR] appears to be a good fit under Ketamine and Fentanyl as the amplitudes of their average impulse responses are

[1]It is sufficient to focus on Z-values only as most fMRI studies base their conclusions about neuronal activity solely on these values, i.e. they are considered to reflect amplitude of neuronal activity without requiring additional qualification.

[2]Despite canonical [ΔHbR] fitting poorly to the average impulse responses under Alpha-Chloralose and Isoflurane, the mean Z values under these two anaesthetics are still relatively large (6.2028 and 5.7716, respectively - see Table 2.4). This is because the amplitudes of the hemodynamics under these two anaesthetics are larger than under any other anaesthetics (more than an order of magnitude larger than the average hemodynamics under Pentobarbital and Propofol).

close the amplitudes of the average canonical [ΔHbR] across most stimulation train durations. However, these anaesthetics do appear to impact mildly on the linearity of hemodynamic responses[1] as there is mismatch between the amplitudes of the average impulse response under these anaesthetics and the GLM response for the 5s, 7s and 9s stimulation train durations, but not for the shorter or longer stimulation train durations, i.e. there are inconsistent differences across the stimulation train durations.

These qualitative observations agree with the Z-value discrepancies calculated in Chapter 3. For example, it is clear from Fig.4.2 that the impulse response under this anaesthetic has a very different shape to that assumed by canonical [ΔHbR]: the timing of the main response of the average impulse response aligns more closely with the timing of canonical [ΔHbR]'s poststimulus overshoot than its main response, and it features a large and prolonged initial dip that overlaps with the main response of canonical [ΔHbR][2]. These peculiarities in impulse response shape are removed for the simulated responses, and so it is unsurprising that the GLM returns a much larger Z-value for these data (the discrepancy is 7.4409 when only canonical [ΔHbR] was included as a basis function in the GLM) .

In the next section, generative modelling is used to separate the contributions of anaesthetic modulations of impulse response shape and hemodynamic linearity on GLM Z-value discrepancies, providing more rigorous qualitative observations to explain the (quantitative) generative embedding results obtained in the second half of this chapter.

[1]The nonlinearity of a hemodynamic response is a function of the stimulation train frequency and duration [81]. This is because hemodynamics experience a cumulative refractoriness when stimuli are delivered in quick succession [116], causing hemodynamics to increase sublinearly with stimulation train duration [80].

[2]It is for these reasons that the negative shape of canonical [ΔHbR] is a better match for the average impulse response under Pentobarbital than canonical [ΔHbR] itself (see Fig.4.2), and why the Z-value calculated for this anaesthetic is negative (see Table 2.2).

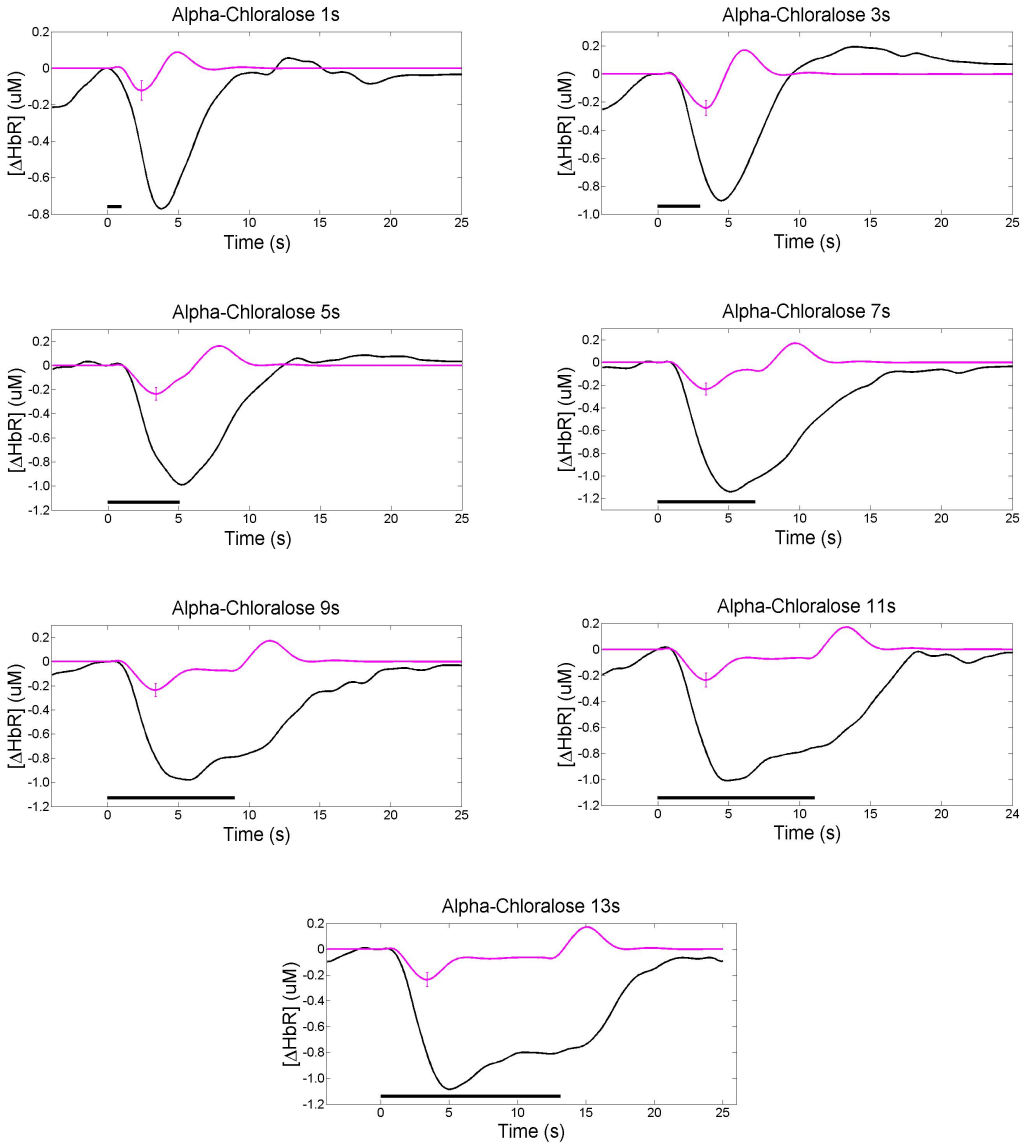

Figure 4.1: Comparison of average GLM responses (magenta) to the real average [ΔHbR] responses (black) under Alpha-Chloralose at each stimulation train duration. The error bar reflects the standard error of the individual canonical [ΔHbR] $\boldsymbol{\beta}$ values.

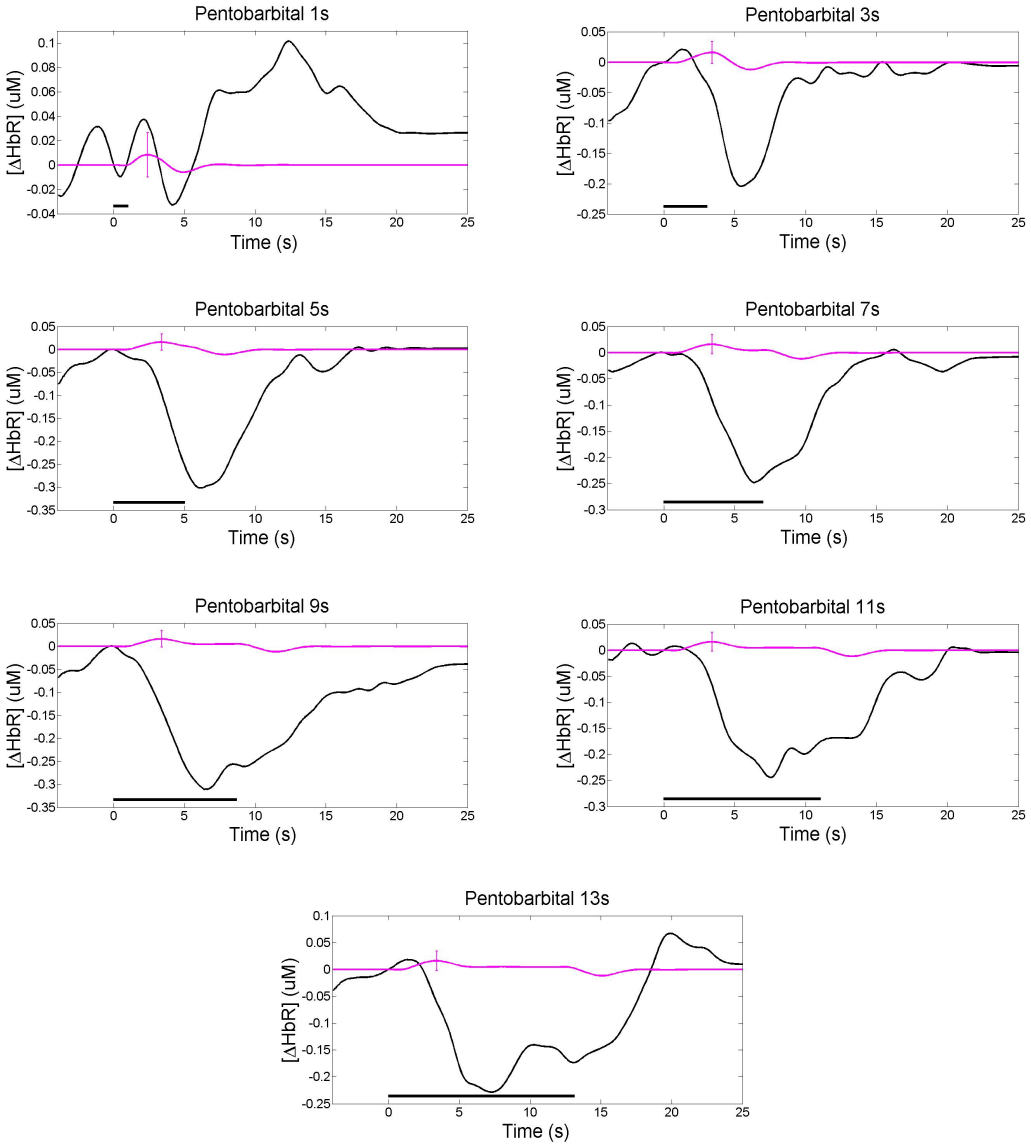

Figure 4.2: Comparison of average GLM responses (magenta) to the real average [ΔHbR] responses black) under Pentobarbital at each stimulation train duration. The error bar reflects the standard error of the individual canonical [ΔHbR] $\boldsymbol{\beta}$ values.

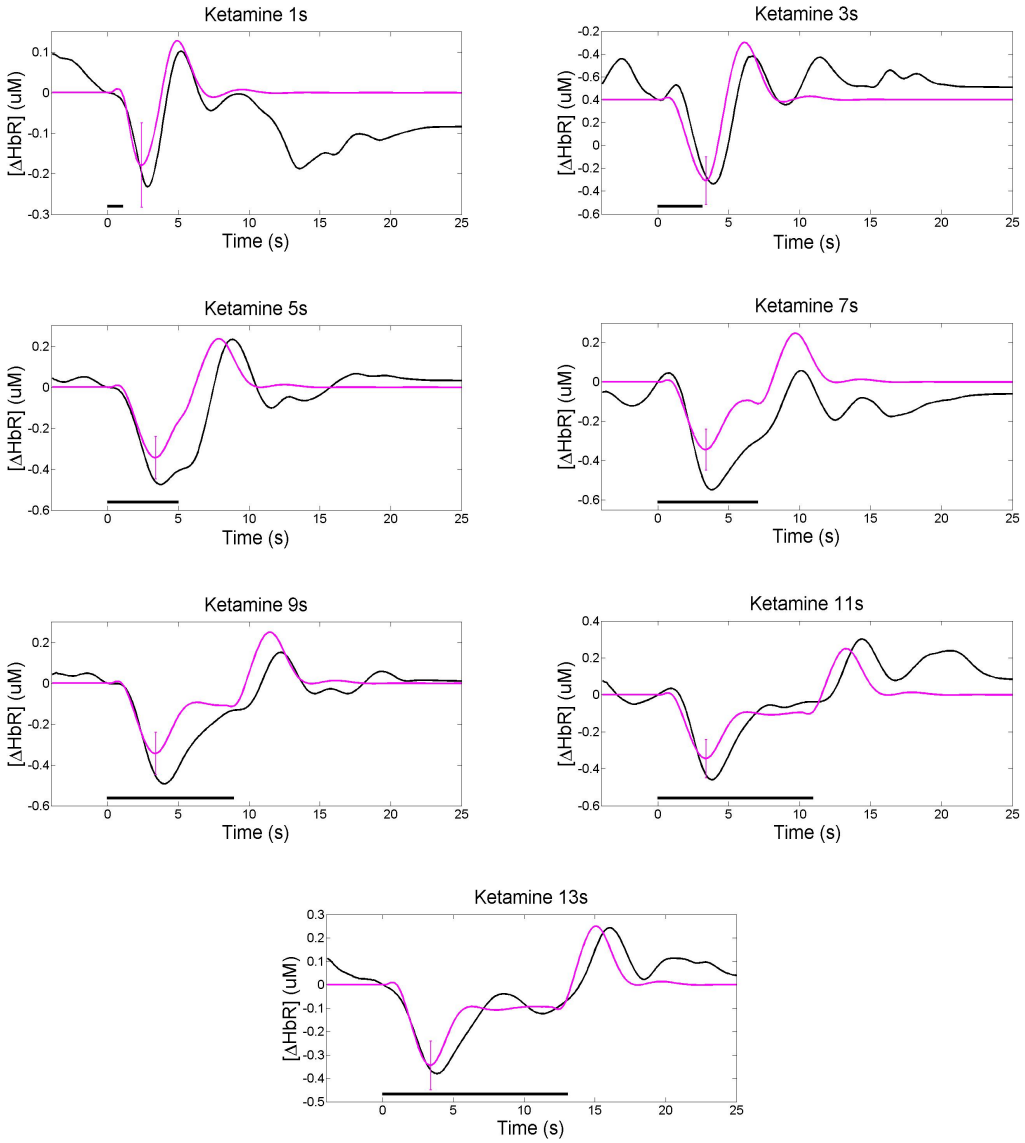

Figure 4.3: Comparison of average GLM responses (magenta) to the real average [ΔHbR] responses (black) under Ketamine at each stimulation train duration. The error bar reflects the standard error of the individual canonical [ΔHbR] $\boldsymbol{\beta}$ values.

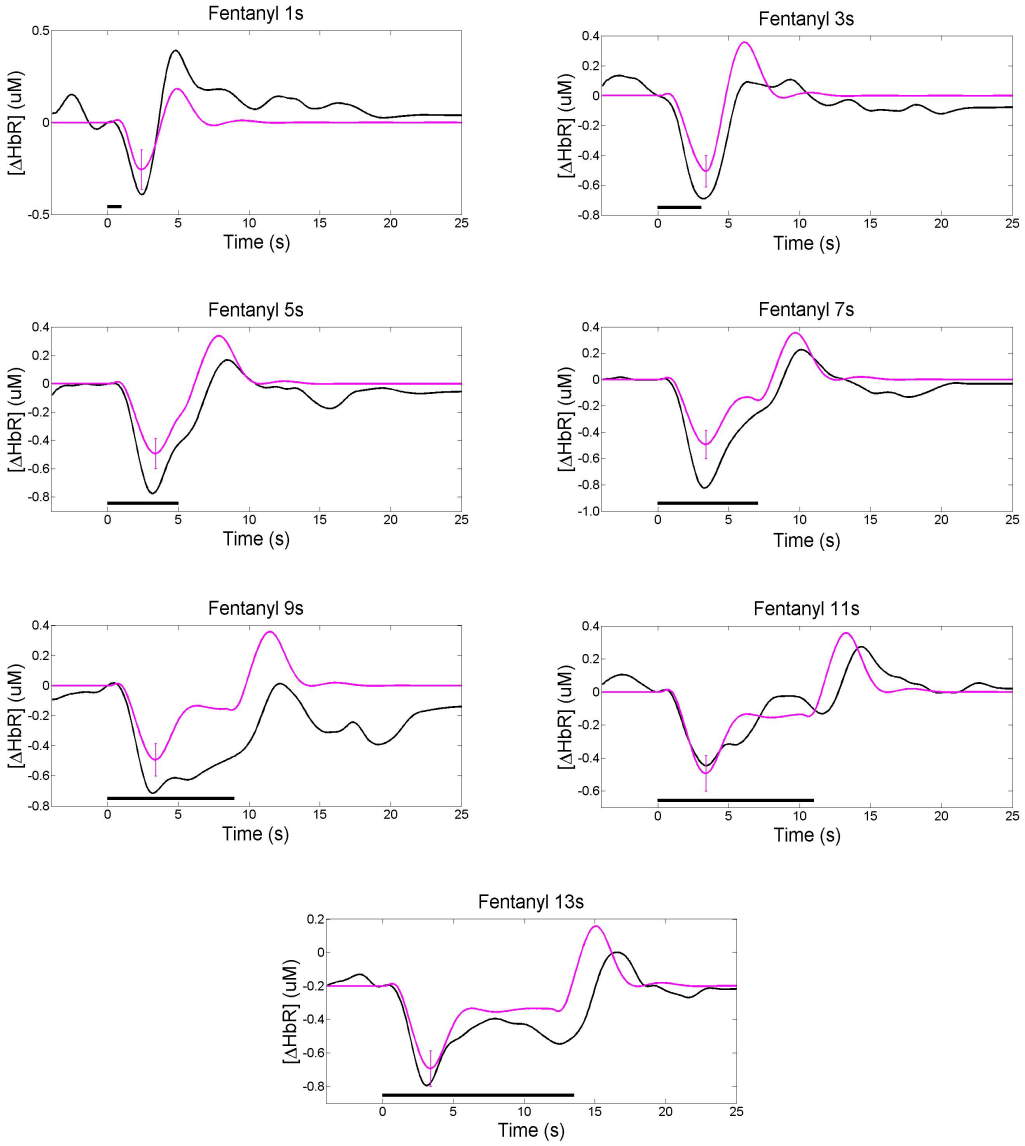

Figure 4.4: Comparison of average GLM responses (magenta) to the real average [ΔHbR] responses (black) under Fentanyl at each stimulation train duration. The error bar reflects the standard error of the individual canonical [ΔHbR] $\boldsymbol{\beta}$ values.

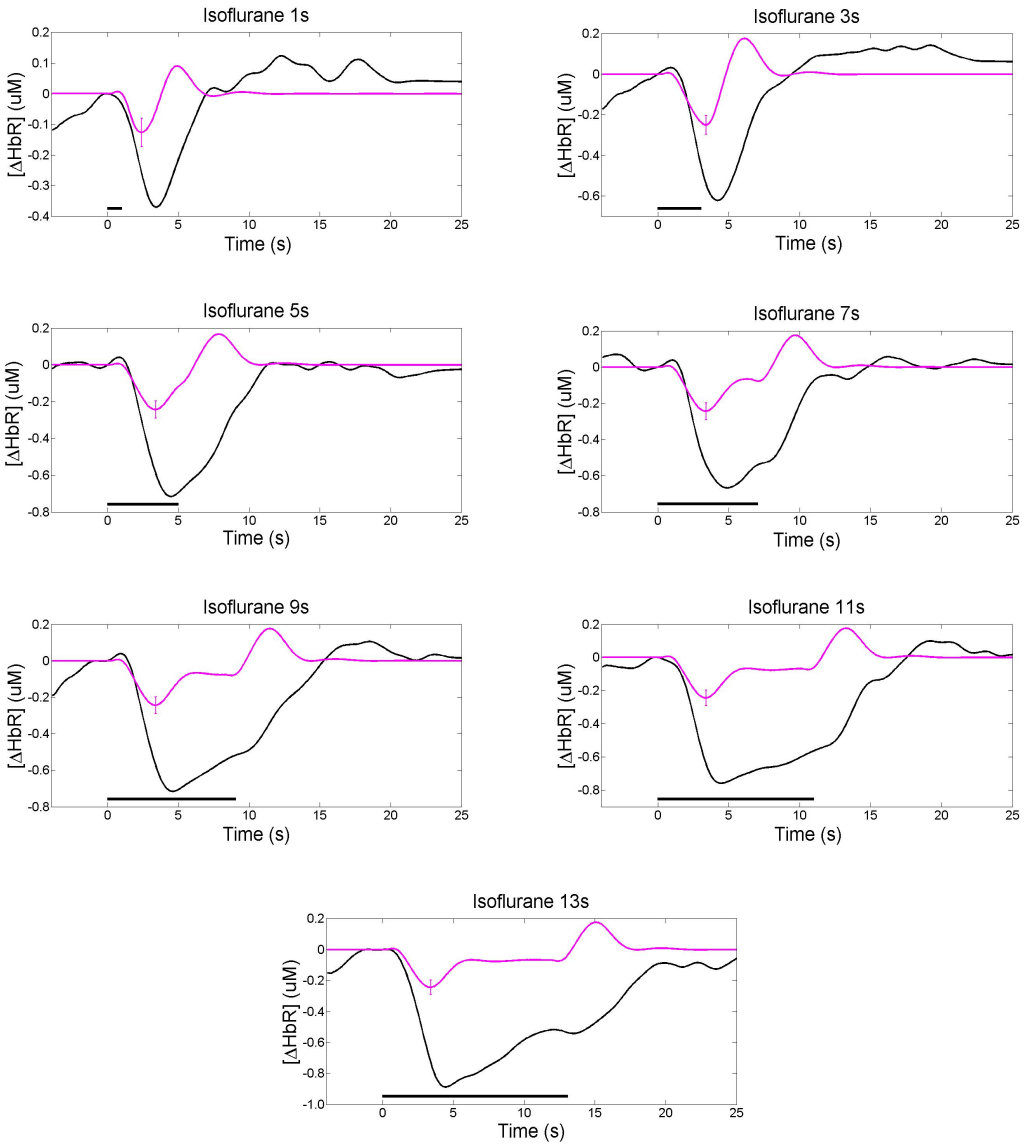

Figure 4.5: Comparison of average GLM responses (magenta) to the real average [ΔHbR] responses (black) under Isoflurane at each stimulation train duration. The error bar reflects the standard error of the individual canonical [ΔHbR] $\boldsymbol{\beta}$ values.

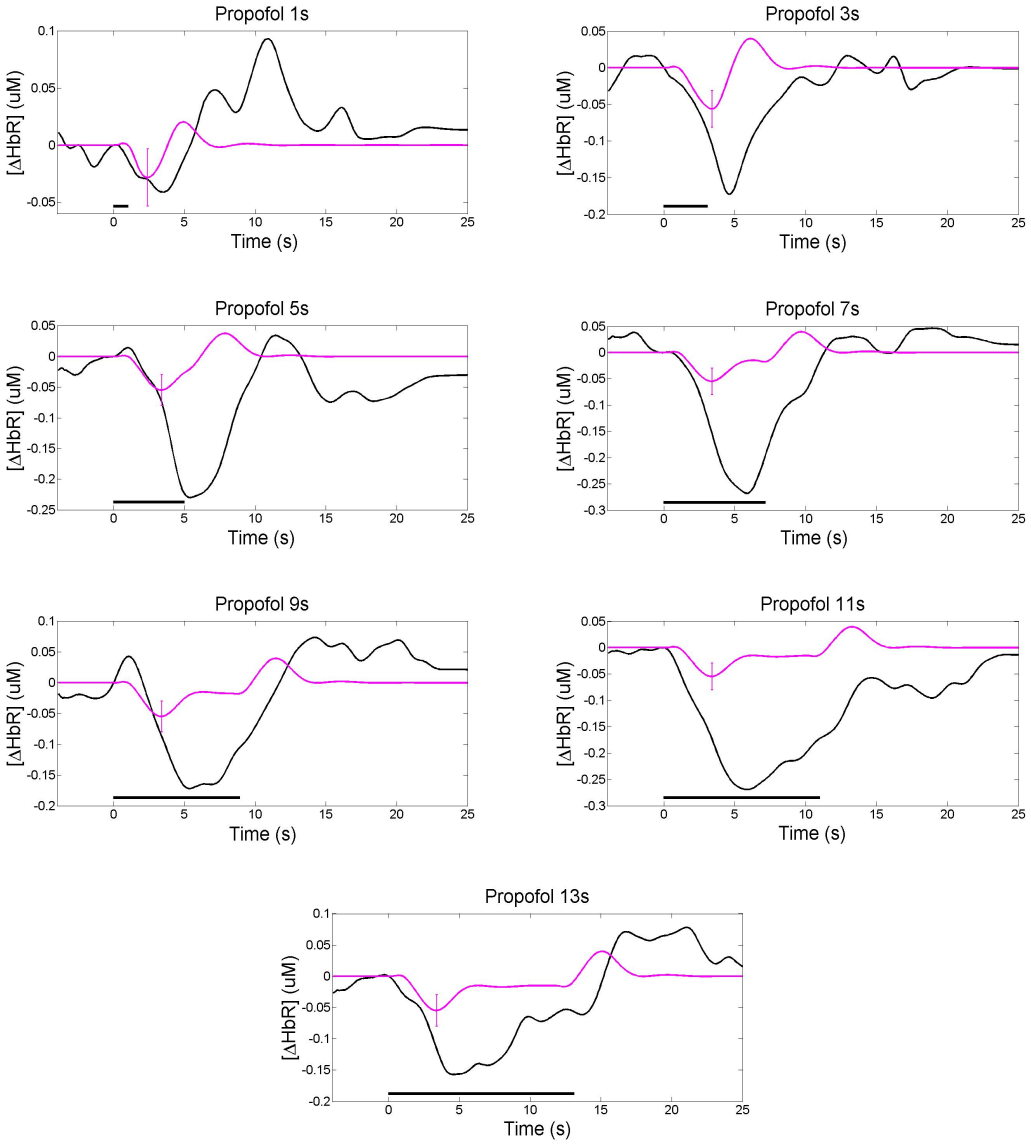

Figure 4.6: Comparison of average GLM responses (magenta) to the real average [ΔHbR] responses (black) under Propofol at each stimulation train duration. The error bar reflects the standard error of the individual canonical [ΔHbR] $\boldsymbol{\beta}$ values.

4.2 Demonstrating how anaesthetics modulate hemodynamic linearity and impulse response shape

Linear neurovascular coupling is a critical assumption of the GLMs used to predict neuronal activities from measures of hemodynamics and any deviations from linearity worsen GLM Z-values. However, it is believed by many that the hemodynamics in healthy brain are only approximately linear [16][104]. Friston et al. showed that the Balloon Model was able to capture the subtle nonlinearities in hemodynamics very well[1] [61] and the small impacts of including these nonlinearities on [ΔHbR] signals from healthy, "unanaesthetised" brain are shown in Fig.4.7.

If neuromodulations such as anaesthetics greatly enhance the nonlinear hemodynamics processes, the assumption of linearity is no longer suitable, resulting in either increased or decreased effect size, and/or increased GLM standard errors. This is also true of neuromodulation of hemodynamic impulse response shape: if the shape of the hemodynamic impulse response is altered by the neuromodulation then the effect size can be either increased or decreased, and/or the standard errors can be increased. As both the effect size and standard error determine individual GLM significances, the consequence of enhanced nonlinearity and/or altered impulse response shape for GLM significance can be either overestimation of neuronal activity (as is the case for Ketamine, Fentanyl and Isoflurane), or underestimation (as is the case for Alpha-Chloralose, Pentobarbital and Propofol).

[1]The Balloon Model captures the following nonlinear phenomena (at least to some extent):

- nonlinear saturation of CBV (so that the Balloon cannot distend indefinitely, which is obviously biologically impossible),
- nonlinear oxygen extraction dependency on CBF,
- nonlinear outflow dependency on CBV.

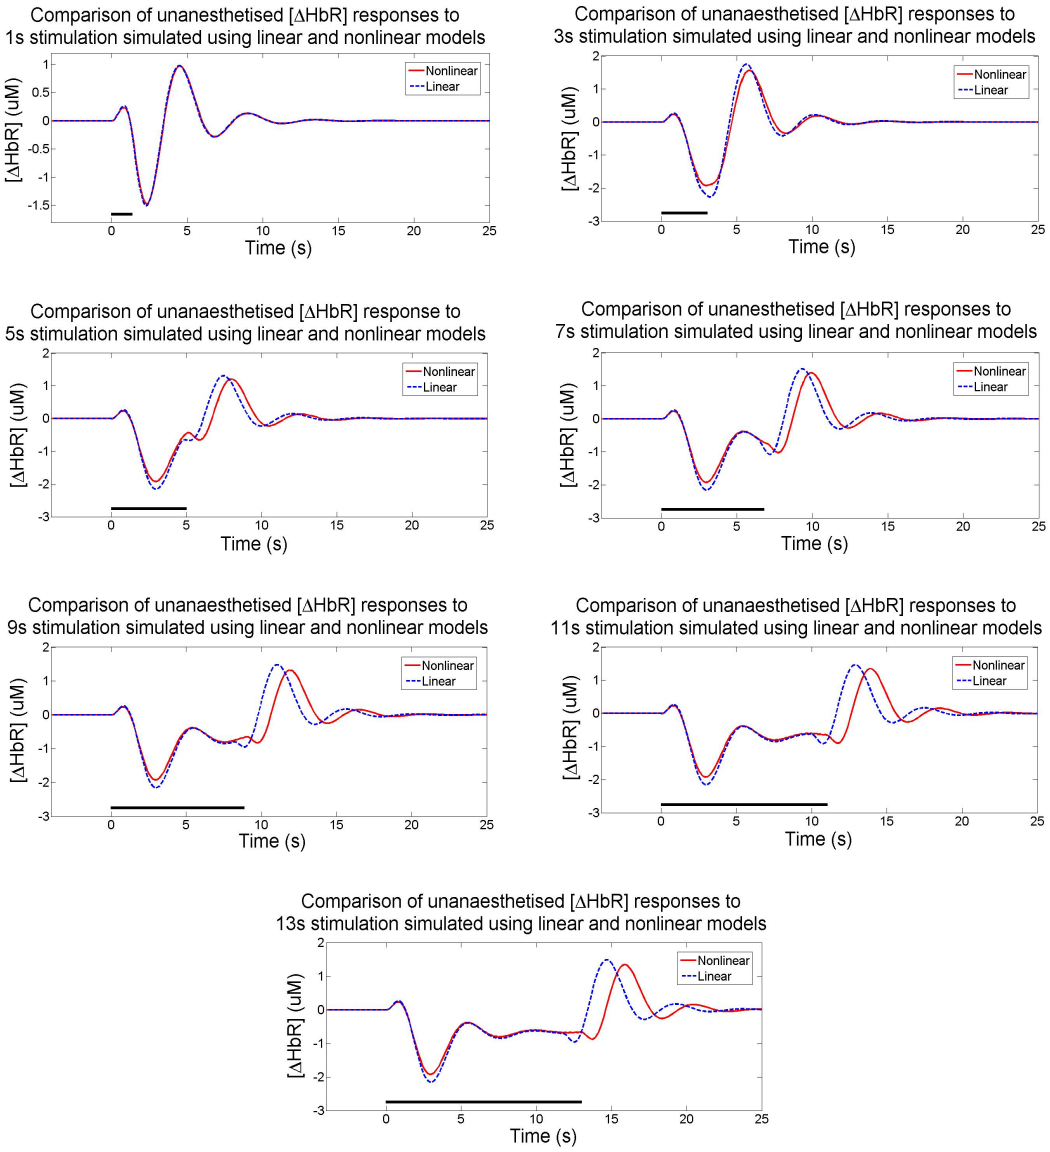

Figure 4.7: Comparison of [ΔHbR] responses simulated by 1) driving Combined DCM (using posterior parameter values from fitting Combined DCM to the awake-rat impulse response function), which contains the nonlinear Balloon Model (red), and 2) linearly convolving the stimuli with the awake-rat impulse response function (blue).

4.2.1 Methods

In order to investigate whether the different anaesthetics modulate hemodynamic linearity and/or impulse response shape, the posterior parameters from Combined DCM fitting to the average 5s stimulation train responses under each anaesthetic were used to simulate the corresponding impulse response functions, and these impulse response functions were used to create: 1) linear responses using linear convolution, and 2) nonlinear responses using Combined DCM, for each stimulation train duration[1]. For the real [ΔHbR] responses under an anaesthetic to be considered to be approximately linear, the simulated responses to all stimulation train durations under that anaesthetic should be similar to the simulated responses generated using linear convolution. If they are not similar, then it can be concluded that that anaesthetic enhances hemodynamic nonlinearities, and that this phenomenon might at least partly explain the Z-value discrepancy that occurred under that anaesthetic.

To instead investigate whether anaesthetic modulations of impulse response shape contribute to the confounding of GLM analyses as well as/instead of anaesthetic modulations of nonlinearities, Combined DCM was fitted to the average response to each stimulation train duration under each anaesthetic, and each of these sets of posterior parameters was then used to simulate a corresponding impulse response[2]. For it to be concluded that an anaesthetic does not modulate impulse response shape, the average impulse response simulated for that anaesthetic (average across the seven stimulation durations) should be similar to canonical [ΔHbR].

[1] It was intended that mDCM-fMRI would be used instead of Combined DCM for this work as it is not necessary for the neuronal processes to be modelled, however, even with multiple parameter prior reinitialisations, the mDCM-fMRI fits to many of the average [ΔHbR] responses were poor.

[2] This idea is analagous to the work carried out by Friston et al., 1998, to validate use of their second-order Volterra kernel series: they estimated the hemodynamic response to hearing a single word from the real responses to hearing multiple words. Because they had the responses to single words available too, they were able to confirm that their second-order Volterra series captured the vast majority of the hemodynamics [62].

4.2.2 Results

Fig.4.8 compares the [ΔHbR] responses simulated for each anaesthetic and stimulation train duration using 1) linear convolution, and 2) (nonlinear) Combined DCM. There are two differences between the linear and nonlinear responses common to all of the anaesthetics:

- The features in the second half of nonlinear responses occur slightly later than they do in the linear responses, meaning that there are clear differences in the timings of their poststimulus overshoots. Importantly, these differences increase with stimulation duration.
- The amplitudes of the linear responses are slightly larger than the amplitudes of the nonlinear responses[1].

The real responses shown in Fig.4.1 - 4.6 are slower than the equivalent (linear) GLM responses, but these differences do not increase with stimulation duration, i.e. they appear to be relatively constant across the stimulation durations, and the dynamics of their first halves are slower than the dynamics of the first half of canonical [ΔHbR] (in addition to their second halves being slower than the second half of canonical [ΔHbR]). It is therefore possible that some of the mismatches of poststimulus undershoot timings seen in the responses to the longer stimulation durations in Fig.4.1 - 4.6 are a consequence of anaesthetic modulations of hemodynamic nonlinearities, but anaesthetics must modulate other aspects of hemodynamics for the remaining discrepancies to be explained.

Fig.4.9 compares the average impulse responses under each anaesthetic with canonical [ΔHbR] (scaled to match the simulated impulse response in amplitude). The main conclusion to make from these plots is that the shapes of the average impulse responses occuring under all of the anaesthetics except for Ketamine or

[1] As GLM amplitude is a free parameter, mismatches in GLM and real response amplitudes cannot be attributed to GLMs not being able to capture nonlinearities, and so this difference between the linear and nonlinear responses is inconsequential to this work.

Fentanyl are considerably different from the shape of canonical [ΔHbR]. Specifically, all of the dynamics in the impulse responses generated under these four anaesthetics are much slower than they are in canonical [ΔHbR]. These differences in the speed of the dynamics are capable of generating the relatively constant mismatches in poststimulus undershoot timings occuring across the stimulation durations as well as the mismatches in the timings of the features in the first halves of the responses. What is more, Fig.4.9 also shows that the relative amplitudes of the initial dips and poststimulus undershoots compared to the amplitudes of the main responses in all of the anaesthetic-specific impulse responses (even under Ketamine and Fentanyl, although to less of an extent) are smaller than they are in canonical [ΔHbR]. Taken together, these anaesthetic modulations of impulse response shape can explain most, if not all, of the discrepancies between the real and GLM responses in Fig.4.1 - 4.6.

4.2.3 Discussion

This chapter has so far shown that the impacts that hemodynamic nonlinearities have on [ΔHbR] responses are limited to slightly slower dynamics in the second half of the responses only. On the other hand, [ΔHbR] responses can be changed dramatically by altering impulse response shape. This indicates that anaesthetic modulations of impulse response shape are more responsible for the mismatches between the GLM and real responses than anaesthetic modulations of hemodynamic nonlinearities. In support of this conclusion is the observation that the impulse responses generated under Ketamine and Fentanyl are most similar to canonical [ΔHbR], and these are also the anaesthetics that caused the smallest Z-value discrepancies.

In the next section, generative embedding is used to identify which biological actions of the anaesthetics cause these various modulations.

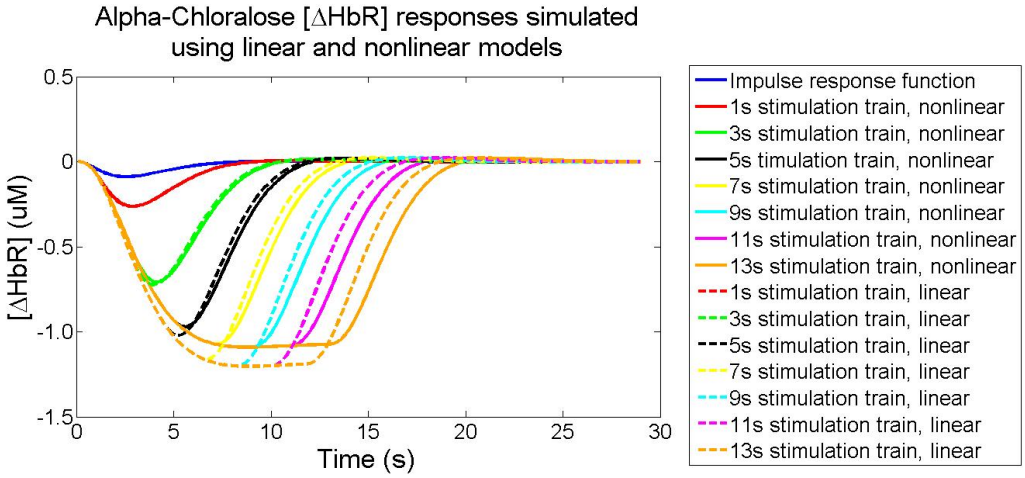
Alpha-Chloralose [ΔHbR] responses simulated
using linear and nonlinear models
[ΔHbR] (uM)
Time (s)
Impulse response function
1s stimulation train, nonlinear
3s stimulation train, nonlinear
5s timulation train, nonlinear
7s stimulation train, nonlinear
9s stimulation train, nonlinear
11s stimulation train, nonlinear
13s stimulation train, nonlinear
1s stimulation train, linear
3s stimulation train, linear
5s stimulation train, linear
7s stimulation train, linear
9s stimulation train, linear
11s stimulation train, linear
13s stimulation train, linear

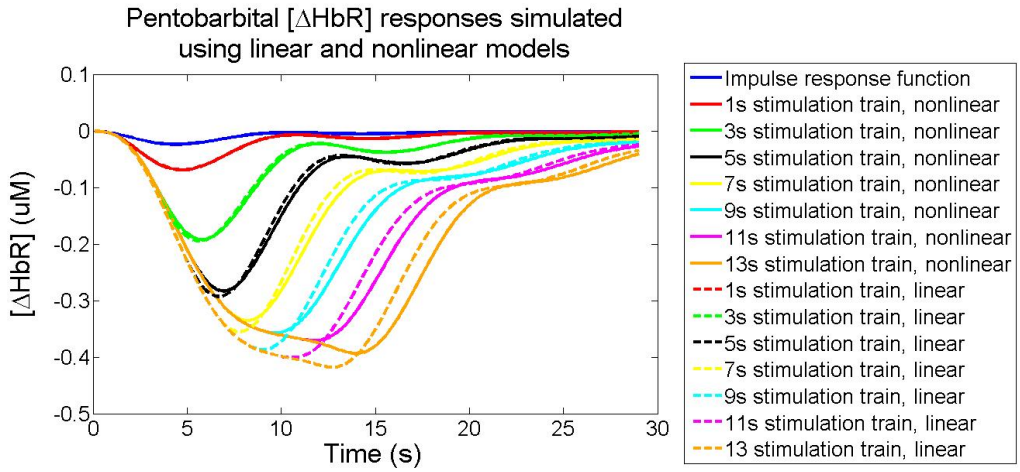
Pentobarbital [ΔHbR] responses simulated
using linear and nonlinear models
[ΔHbR] (uM)
Time (s)
Impulse response function
1s stimulation train, nonlinear
3s stimulation train, nonlinear
5s stimulation train, nonlinear
7s stimulation train, nonlinear
9s stimulation train, nonlinear
11s stimulation train, nonlinear
13s stimulation train, nonlinear
1s stimulation train, linear
3s stimulation train, linear
5s stimulation train, linear
7s stimulation train, linear
9s stimulation train, linear
11s stimulation train, linear
13 stimulation train, linear

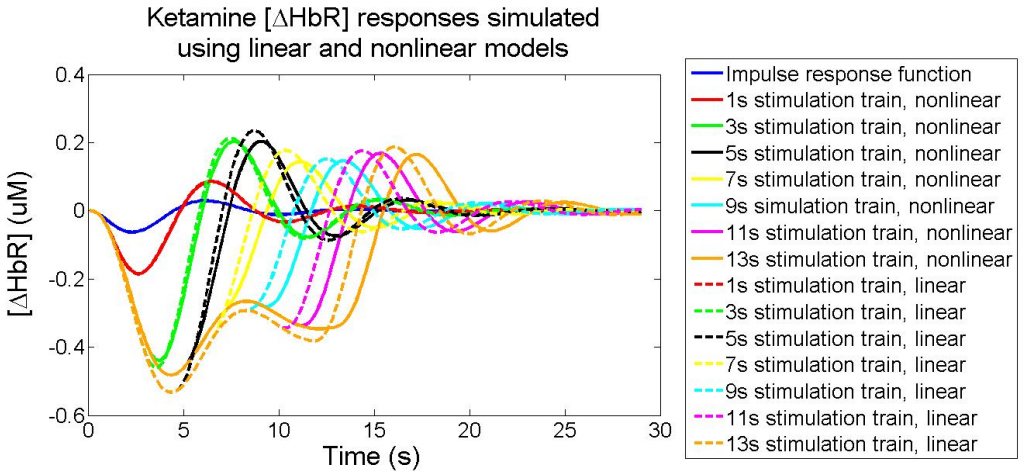
Ketamine [ΔHbR] responses simulated
using linear and nonlinear models
[ΔHbR] (uM)
Time (s)
Impulse response function
1s stimulation train, nonlinear
3s stimulation train, nonlinear
5s stimulation train, nonlinear
7s stimulation train, nonlinear
9s simulation train, nonlinear
11s stimulation train, nonlinear
13s stimulation train, nonlinear
1s stimulation train, linear
3s stimulation train, linear
5s stimulation train, linear
7s stimulation train, linear
9s stimulation train, linear
11s stimulation train, linear
13s stimulation train, linear

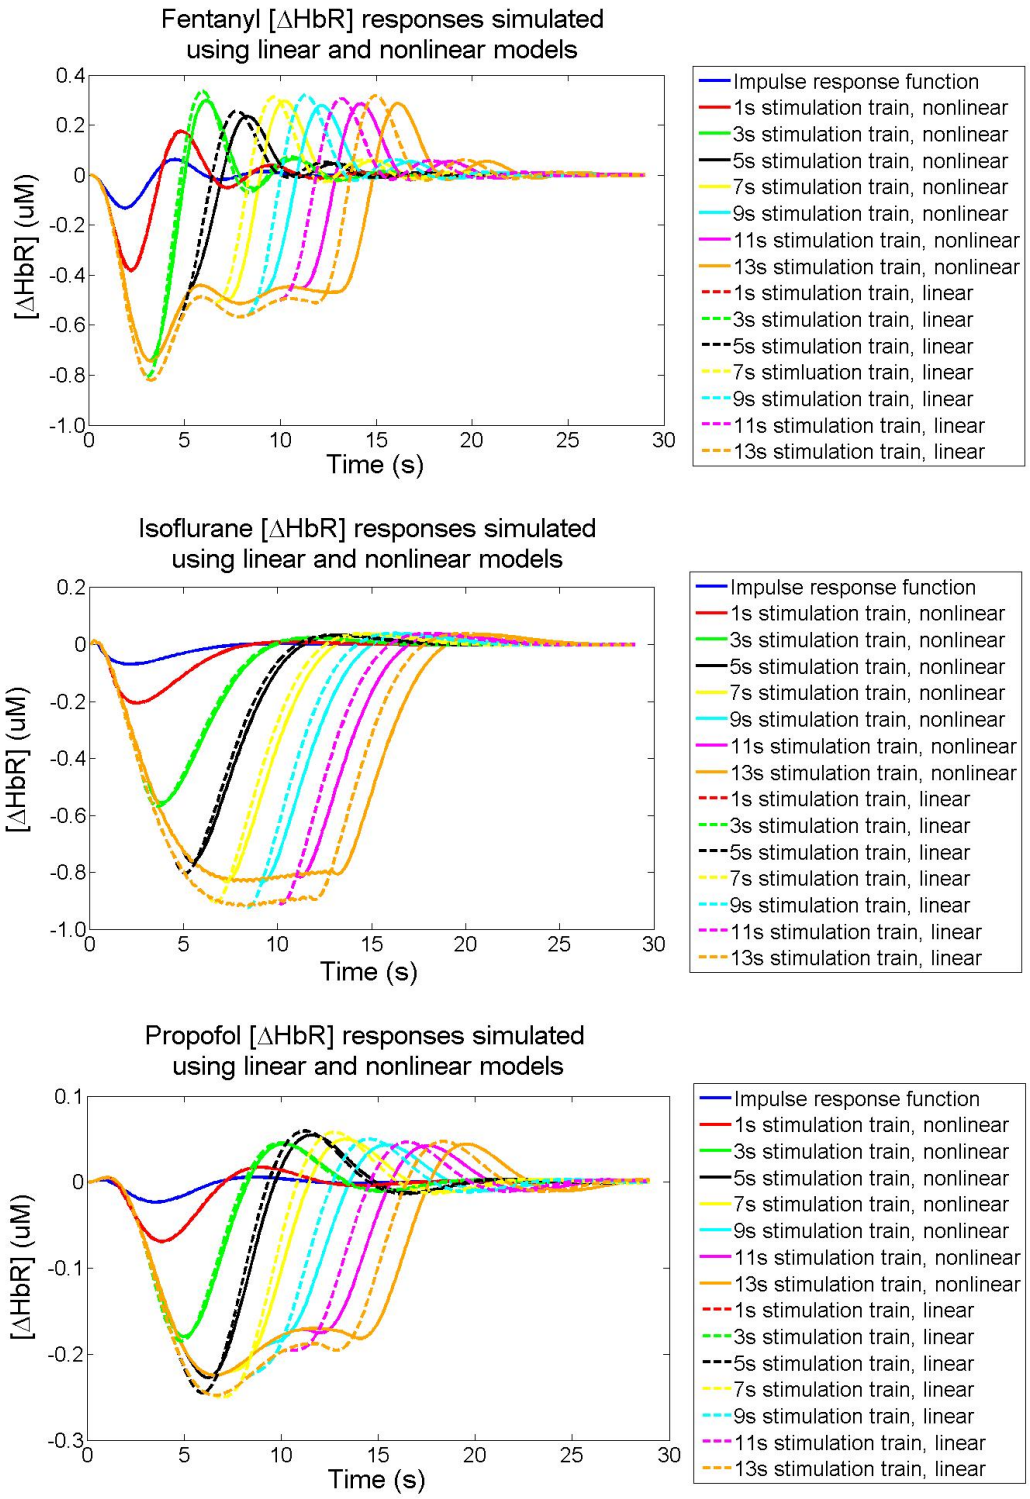

Figure 4.8: Comparison of [ΔHbR] responses simulated using either linear convolution (as is done for GLM) or (nonlinear) Combined DCM for each anaesthetic and stimulation train duration.

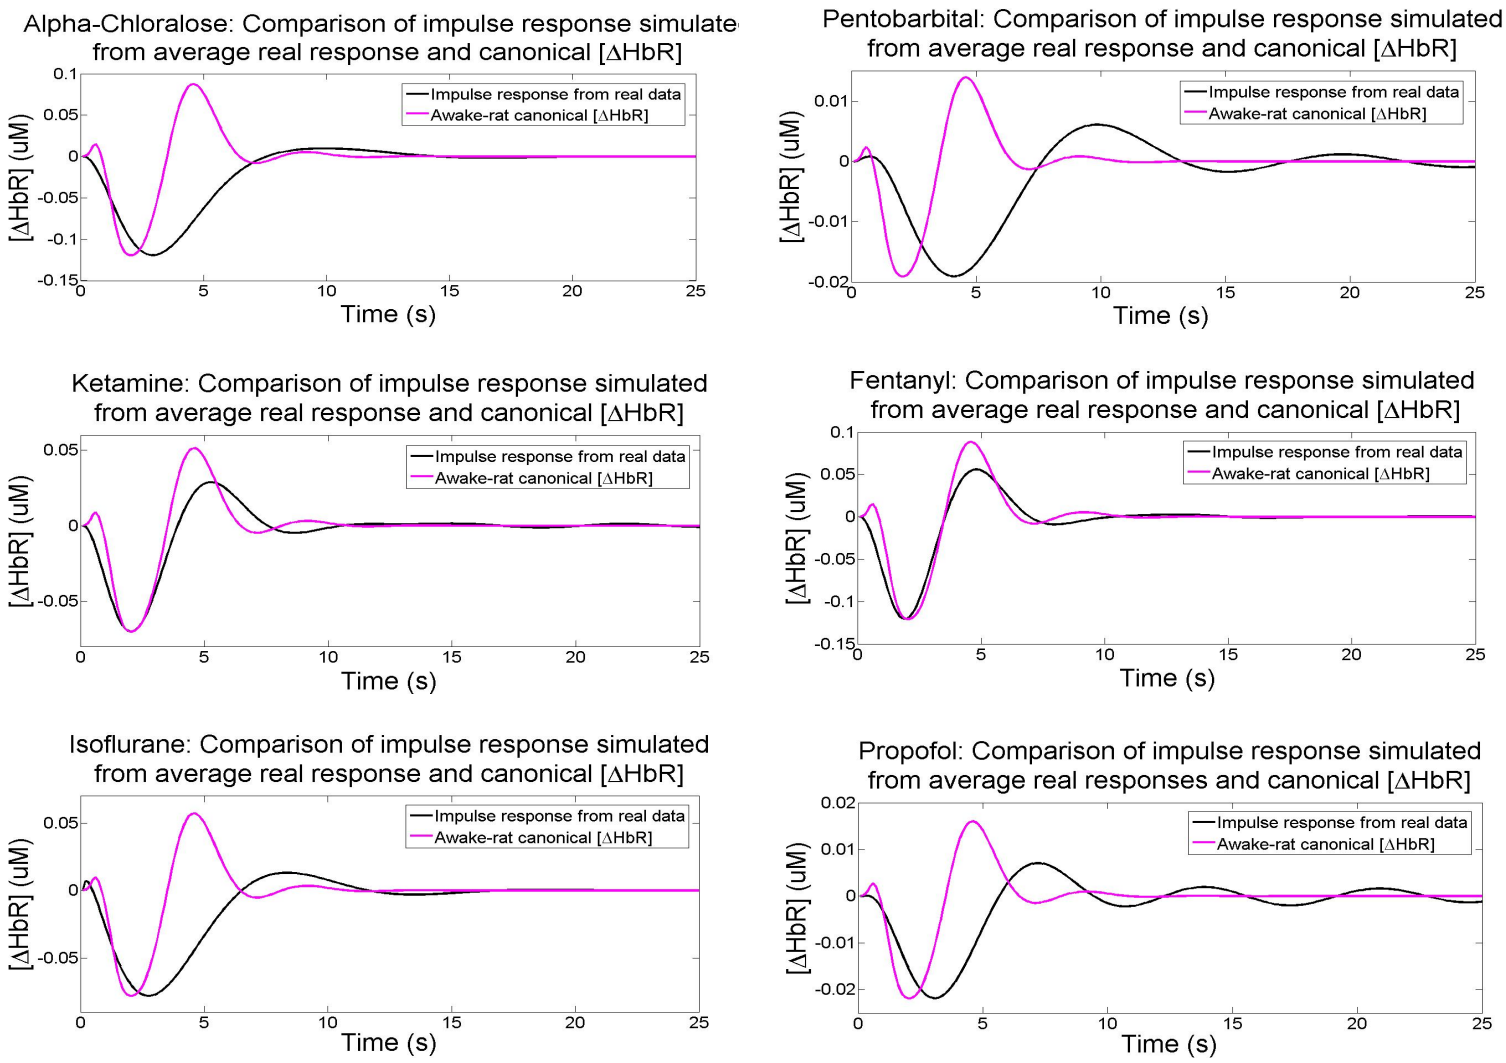

Figure 4.9: Comparisons of 1) the average impulse responses (average across stimulation train durations) simulated for each anaesthetic (black), and 2) canonical [ΔHbR] (magenta), scaled to have identical response amplitudes.

4.3 Understanding which specific biological actions of anaesthetics confound GLM analyses

The aim of this section is to identify which specific biological actions of the anaesthetics cause the modulations of impulse response shape (and possibly the enhancement of hemodynamic nonlinearities) that result in GLMs being confounded to different extents. Because of the small number of Z-value discrepancies available (on average only ten per anaesthetic), this was achieved using generative embedding incorporating simple linear regression[1].

4.3.1 Methods

The following linear regression model (Equ.4.1) was fitted for each of the four different types of GLM (1) canonical [ΔHbR] only, 2) with temporal derivative, 3) with dispersion derivative, and 4) with both derivatives) under each of the six anaesthetics in a stepwise fashion, i.e. by adding each standardised Combined DCM posterior parameter to the model one by one and then using the significance of the improvement in prediction of the Z-value discrepancies to assess whether to keep that parameter as a regressor in the model (or to set its β value to 0)[2]:

$$\mathbf{Z_{real}} - \mathbf{Z_{sim}} = \sum_{i=1}^{5} \sum_{j=1}^{27} \beta_{ij} A_j \, \boldsymbol{p}_i. \tag{4.1}$$

The regressors of this model found to significantly explain the Z-values

[1] If more datapoints were available, a more complete analysis, such as one based on information theory, which is sensitive to all orders of correlations rather than only linear relationships, could have been carried out instead. As mentioned previously, it was not possible to generate more datapoints because responses averaged over fewer repeats were too noisy to reasonably fit Combined DCM to.

[2] Ideally, the uncertainties in the Combined DCM fits would be passed through these linear regression models so that the conclusions made about the influences of the parameters on GLM significances would have error bars. However, the complexity that this would add to the linear regression models makes this an inappropriate calculation to make on this small dataset.

calculated under each anaesthetic reflect the biological processes that, when modulated by anaesthesia, alter the impulse response shape (and possibly hemodynamic nonlinearities) that impact on GLM significance.

The hemodynamics parameters found to significantly explain the Z-value discrepancies under each anaesthetic reflect the hemodynamic processes that directly confound GLM interpretation when modulated by anaesthesia[1]. However, because the neuronal activity driving both the real and simulated hemodynamics should at least be very similar[2], significantly influential DCM-EEG parameters reflect neural processes that confound GLM significance indirectly by altering the vulnerability of hemodynamics processes to anaesthetic modulation. In other words, the $\mathbf{Cp_{EEG-fMRI}}$ between significantly influential DCM-EEG parameters and some/all of the Sotero/mDCM-fMRI parameters alter the extent to which the hemodynamics confound GLMs. For example, if DCM-EEG parameter C (which controls the magnitude of the thalamic input to the cortical column) was found to significantly explain Z-value discrepancies, this would indicate that the modulation of the coupling of C to the hemodynamics by the anaesthetic condition was significantly confounding GLM Z-values under that anaesthetic, i.e. a low C value might mean that anaesthetic condition has little influence on the hemodynamics of relevance to GLM Z-value, but a high C value might increase this vulnerability, or vice versa.

[1] Although response amplitude, onset time and width ought to be captured by GLMs through canonical ΔHbR $\boldsymbol{\beta}$ values, the temporal derivative and the dispersion derivative, respectively, the DCM parameters that contribute to these response features have been included in the linear regressions. This is because if inclusion of the derivatives in the GLM does alter Z-values by better capturing the response onset time and width (and in doing so the amplitude too), Z-value discrepancies ought to be more significantly explained by these parameters when the derivatives are not included in the GLM than when they are.

[2] There may be small differences in the neuronal activities driving the real and simulated hemodynamics, resulting from DCM-EEG not fitting to the real EEG responses with 100% accuracy.

4.3.2 Results and discussion

The Combined DCM parameters that significantly explain the differences in the Z-values calculated for the real and simulated [ΔHbR] signals under each anaesthetic are listed in Tables 4.1 - 4.6. Plots highlighting their linear relationships with Z-value discrepancy are shown in Fig.4.10.

There are two particularly noteworthy observations to make from these results:

- The majority of Combined DCM parameters found to explain GLM Z-value discrepancies across the six anaesthetics are neuronal as opposed to metabolic or hemodynamic. For example, when only canonical [ΔHbR] is used as a basis function in the GLM, all of the Z-value discrepancies were significantly explained by anaesthetic modulations of neuronal processes except for Alpha-Choralose, which was explained by hemodynamic process γ (rate of bloodflow autoregulatory feedback), and Isoflurane, which was explained by metabolic processes a_i and c. This suggests that GLMs are confounded by the indirect effects of anaesthetics on hemodynamics, i.e. through the influences of anaesthetic-modulated neuronal activity on hemodynamics, more than by their direct effects on hemodynamics.
- There are some differences in which Combined DCM parameters significantly explain Z-value discrepancies across the different GLM types, highlighting that different GLM types are sensitive to different aspects of hemodynamics. For example, Combined DCM parameter γ significantly explains the Z-value discrepancies for each GLM type under Alpha-Chloralose except when both temporal and dispersion derivatives are used as basis functions in the GLM, when $\gamma 3$ (connection strength between pyramidal cells and inhibitory interneurons) significantly explains them instead. When only canonical [ΔHbR] and its temporal derivative are included as basis functions in the GLM, $\gamma 1$ (connection strength between pyramidal cells and stellate cells) significantly explains them in addition to γ. These results indicate that when

the temporal derivative is included as a basis function in the GLM[1], i.e. when the onset times of Alpha-Chloralose responses are more accurately captured by a GLM, the actions of the anaesthetic on the connection strengths between the neuronal cell populations (or, rather, the impacts that these connection strengths have on hemodynamics processes) can be detected.

4.4 Overall discussion

As the GLM significance rates and levels calculated for the simulated data with unanaesthetised hemodynamics in Chapter 3 were at odds with those calculated for the real data in Chapter 2, it is clear that anaesthetics are problematic interest neuromodulations for BOLD experiments. The aim of this chapter was to gain insight into which specific actions of the anaesthetics cause them to confound GLM analyses.

In the first half of this chapter, it was shown that hemodynamic nonlinearities cannot be responsible for the majority of the mismatch between the GLM and real responses, whereas anaesthetic modulations of impulse response shape can be. As these are the only ways in which neuromodulations can impact on GLM significances, it was concluded that anaesthetic modulations of impulse response shape are most strongly responsible for confounding GLM analyses of hemodynamics data, explaining why the smallest Z-value discrepancies were recorded under Fentanyl and Ketamine.

In the second half of this chapter, Combined DCM parameters that significantly explain the hemodynamic modulations of consequence to GLM significance were identified.

By combining these results with those made in Chapter 2, the following main insights have been made about the usage of anaesthetics as interest

[1]Notice the small discrepancy in onset times between the average Alpha-Chloralose response and canonical [ΔHbR] in Fig.2.6.

GLM	Canonical [ΔHbR] only	With temporal derivative		With dispersion derivative	With both derivatives
Combined DCM parameter	γ	$\gamma 1$	γ	γ	$\gamma 3$
Regression coefficient	-5.2126	5.0612	-5.2793	-5.1251	-4.9843

Table 4.1: Combined DCM parameters that significantly explain Z-value discrepancies under Alpha-Chloralose.

GLM	Canonical [ΔHbR] only		With temporal derivative		With dispersion derivative		With both derivatives	
Combined DCM parameter	T_e	H_e	T_e	H_e	T_e	H_e	T_e	H_e
Regression coefficient	2.7786	-8.6198	2.9178	-8.6517	2.7316	-8.5485	2.7230	-8.6470

Table 4.2: Combined DCM parameters that significantly explain Z-value discrepancies under Pentobarbital.

GLM	Canonical [ΔHbR] only		With temporal derivative	With dispersion derivative	With both derivatives
Combined DCM parameter	$\gamma 1$	$\gamma 2$	$\gamma 2$	$\gamma 2$	$\gamma 2$
Regression coefficient	4.6354	-19.0543	-19.4113	-19.0419	-19.5987

Table 4.3: Combined DCM parameters that significantly explain Z-value discrepancies under Ketamine.

GLM	Canonical [ΔHbR] only			With temporal derivative				With dispersion derivative			With both derivatives		
Combined DCM parameter	T_i	H_i	C	T_i	H_i	C	c	T_i	H_i	C	T_i	H_i	C
Regression coefficient	1.8927	5.5087	-6.7885	1.8989	5.4420	-6.7363	3.8393	1.8072	5.2214	-6.4477	1.8211	5.4152	-6.5440

Table 4.4: Combined DCM parameters that significantly explain Z-value discrepancies under Fentanyl.

GLM	Canonical [ΔHbR] only			With temporal derivative	
Combined DCM parameter	$\gamma 2$	a_i	c	R2	γ
Regression coefficient	-10.9468	-2.7554	-0.0379	-10.4750	-4.9504

With dispersion derivative								With both derivatives	
H_e	S1	$\gamma 1$	$\gamma 2$	$\gamma 3$	t_e	t_i	c	$\gamma 2$	c
-10.1052	-7.7222	3.9131	-10.5367	7.5041	-5.4631	-1.0236	0.0335	-10.2346	0.0339

Table 4.5: Combined DCM parameters that significantly explain Z-value discrepancies under Isoflurane.

GLM	Canonical [ΔHbR] only	With temporal derivative		With dispersion derivative		With both derivatives
Combined DCM parameter	$\gamma 4$	$\gamma 4$	τ	$\gamma 4$	τ	$\gamma 4$
Regression coefficient	7.5940	7.9683	1.7451	7.8867	1.7693	7.9712

Table 4.6: Combined DCM parameters that significantly explain Z-value discrepancies under Propofol.

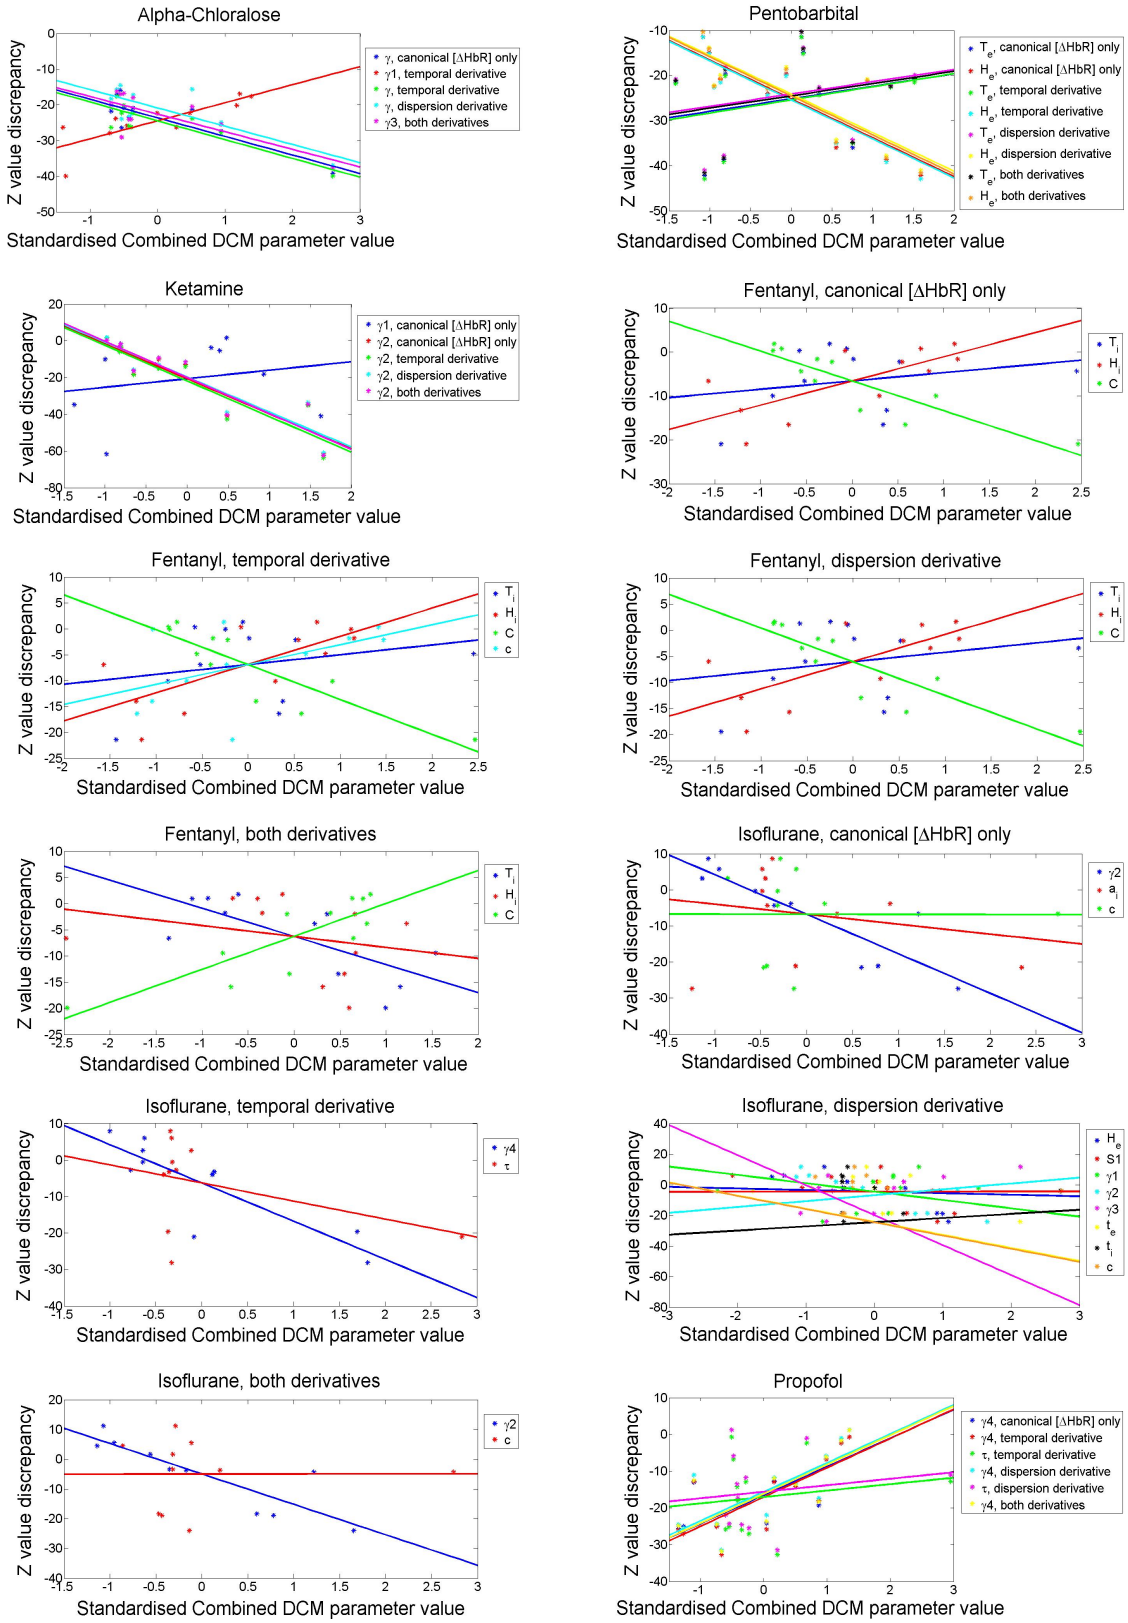

Figure 4.10: Combined DCM parameters plotted against the discrepancies in GLM Z-values for the real and simulated [ΔHbR] signals that they significantly (linearly) explain. When more than one or two Combined DCM parameters significantly explain Z-value discrepancies (i.e. under Fentanyl and Isoflurane), these are shown separately for each GLM model type.

neuromodulations in BOLD experiments:

- Of the six anaesthetics investigated in this thesis, Fentanyl is the recommended neuromodulation if only the standard BOLD GLM (with or without canonical HRF derivative(s) as basis function(s)) is used to detect neuronal activations. This is because Fentanyl confounds GLM analyses the least, allowing neuronal activity to be detected with correct significance levels.
- If it is necessary to use one of the anaesthetics that confound standard GLM analyses, one or more of the algorithms described at the end of Chapter 3 could be added to the standard GLM to reduce inaccuracies in the conclusions made about neuronal activity, informed by the nonlinearity and impulse response shape findings made using generative modelling in the first half of this chapter and/or the generative embedding results from the second half of this chapter.

Having practical methods for improving the accuracy of BOLD studies generally should help to increase adoption of fMRI outside academia, e.g. in clinical settings and by drug manufacturers. In the following chapters of this thesis, the potential of generative embedding algorithms for applications specific to these industries, namely for diagnosing, monitoring and understanding neurodegenerative diseases, is explored.

Chapter 5

Do neurodegenerative diseases alter neurovascular coupling?

This thesis has so far focussed on the challenges that neuromodulations pose for BOLD-fMRI because of their effects on neurovascular coupling. The consequence of this work has great significance for the current use of BOLD-fMRI, which is more or less limited to comparing the brain function[1] of different patient groups in cognitive and clinical psychology research. However, the potential of BOLD-fMRI signals extends to widespread clinical use in hospitals for routine diagnosis and monitoring of patients with brain diseases and disorders, as well as to drug development. The remainder of this thesis concentrates on evaluating whether the effects of disease neuromodulations on neurovascular coupling, which so far have only been viewed as problematic for BOLD-fMRI, can instead be exploited for these applications.

In this chapter, a neurodegenerative-disease preliminary study analysed using the traditional analysis approach (linear regression of shape features) is presented in order to provide some understanding of the power with which a prototype hybrid neuroimaging system created by Dartmouth Brain Imaging Group

[1]The emerging trend is to investigate the activation/deactivation of brain networks rather than individual brain regions.

(Hanover, New Hampshire, USA) can discriminate healthy and diseased brains. This work is preparatory for Chapter 6, which aims to:

1. Quantify the discriminability of these disease states if the neuroimaging signals were analysed using generative embedding rather than the traditional analysis method.
2. Calculate the increase in disease state discriminability that occurs if the neuronal activity and hemodynamics are imaged concurrently rather than separately.

5.1 BOLD-fMRI as a tool for diagnosing/monitoring/developing drugs for brain diseases

Diagnosing brain diseases using neuroimaging signals is a method of decoding. Until recently, decoding algorithms for neuroimaging signals have focussed on classifying brain states using the macroscopic spatiotemporal patterns of brain activations predicted from BOLD-fMRI signals [120][73], providing a way for the spatiotemporal brain activation patterns of patients to be compared with disease-specific templates in order to calculate probabilities for them to be suffering from different neural diseases. This clinical investigation would require the collection and then analysis of a large 4D volume (3D brain over time) of BOLD-fMRI signals, and so advancements need to be made to speed it up and reduce its cost for it to become a viable routine clinical practice.

Dartmouth Brain Imaging Group has developed prototypes of an alternative system for detecting brain diseases based on neurovascular coupling that is fast and cheap enough to be routinely used in clinical practice [67][1]. The speed and

[1]Dartmouth Brain Imaging Group's vision is for their product to become the "ECG for brain health" [53]

relative cheap cost of this system is enabled by the fact that neurovascular coupling can be estimated from signals of concurrent neuronal activity and hemodynamics from discrete brain regions over time, i.e. from a set of 2D spaces.

Currently, Dartmouth Brain Imaging Group's system relies on software based on the traditional analysis (linear classification of the multimodal shape parameter space) for classifying brain disease state [131]. Given that the superiority of generative embedding algorithms over the traditional analysis for investigating hypotheses related to neurovascular coupling has already been demonstrated in the first half of this thesis, the aim of the remainder of this thesis is to understand the potential of generative embedding for improving brain health by interrogating aspects of neurovascular coupling.

Brodersen et al., 2010 and 2011, have already gone some way to demonstrate the utility of generative embedding for diagnosing brain diseases: they fitted DCM-LFP to LFP signals collected from a single brain region in rats under two conditions and showed that it was possible to classify the responses by their DCM posterior means with significantly above-chance accuracy in most cases (mean accuracy of 83.6% for distinguishing which of two whiskers had been stimulated) [18]. They also classified aphasic patients from controls using the (macroscopic) DCM-fMRI connection parameters between brain regions estimated from BOLD signals with 98% accuracy [19]. Given the high rates of classification accuracy achievable using generative embedding, as well as the advantages of generative embedding analyses over analyses that require patterns of brain function to be estimated, it appears that the prospects of generative embedding for diagnosing brain diseases are good.

Improving brain health is not limited to diagnosing brain diseases (using classification algorithms) because the biological interpretability afforded by generative embedding analyses could extend the applications of any hardware developed to measure concurrent signals of neuronal activity and hemodynamics to

also monitoring brain diseases or informing drug development (using, for example, ANOVA statistics to detect significant differences in specific biological aspects of neurovascular coupling between timepoints). All of these applications would benefit from increased discriminability of the healthy and diseased feature spaces, i.e. multidimensional spaces of shape features/generative parameters that are more easily separated into healthy/disease classes should allow for healthy/disease subjects to be classified more accurately and for their disease progression/improvement to be detected with greater statistical significance. This discriminability may be enhanced if information about the coupling of the neuronal activity and hemodynamics signals is available.

5.1.1 Neurodegenerative diseases

Neurodegenerative diseases are a group of diseases characterised by impairment and death of neurons. Sub-types include Alzheimer's disease, Parkinson's disease and Prion disease. Although these sub-type diseases suffer some distinctive symptoms, many of the sub-cellular processes involved in their development are similar, meaning not only that common therapies might exist, but also that there may be common detection and diagnosis strategies.

Neurodegenerative diseases are currently most often diagnosed using neuropsychological testing and structural MRI once cognitive decline has already occurred (as that is when symptoms become obvious) [57]. Yet in order to provide the most effective intervention and care program possible, diagnosis needs to be made in the earliest stages of disease [156]. There is evidence that neurodegenerative changes in the brain can precede noticeable cognitive decline by decades, and work is therefore being done to realise what the earliest biomarkers might be. For Alzheimer's disease, these efforts include detection of abnormal amyloid-β (Aβ) concentrations or tau-mediated neuronal injury and dysfunction using either positron-emission tomography [87] or cerebrospinal fluid testing [68].

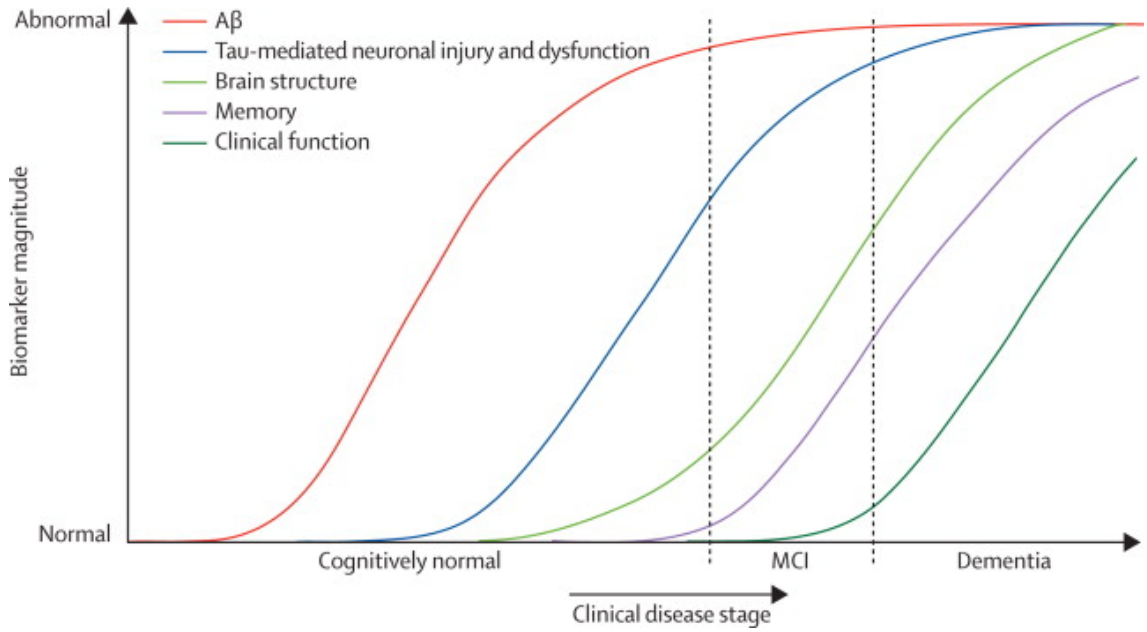

Figure 5.1: Model of Alzheimer's disease-specific early biomarkers. Current disease diagnosis using structural imaging is only possible when brain structure is sufficiently abnormal. Figure reproduced from Jack et al., 2013, Figure 1 [83].

Fig.5.1 models the progression of these biomarkers relative to disease progression measures.

This work proposes an alternative biomarker of neurodegenerative disease, namely modulation of neurovascular coupling, which may be effective even earlier in disease progression than the disease-specific biomarkers highlighted in Fig.5.2.

The processes that lead to cerebrovascular dysfunction in early-stage neurodegenerative disease are hypothesised to be inflammation → glial cell effects → cerebrovascular dysfunction. This is because 1) central inflammation accompanies neurodegenerative diseases [150], 2) microglia (macrophages of the central nervous system) are in an activated state during central inflammation [150], and therefore 3) cerebrovascular regulation, which is maintained by the "functional unit" comprising neurons, cerebral blood vessels and, importantly, glial cells such as astrocytes (the neural cells tasked with supporting and maintaining neuronal health - see Fig.5.3), becomes impaired in early-stage neurodegenerative disease [82].

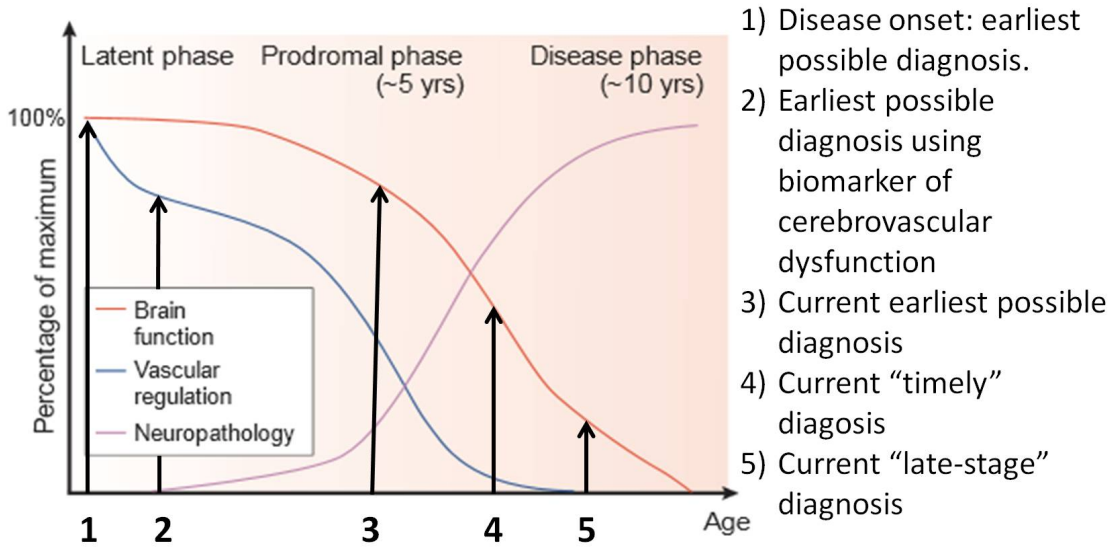

Figure 5.2: Theoretical timeline of neurodegenerative disease progression. Figure reproduced from Iadecola, 2004, Figure 7 [82].

It is unclear what the role of cerebrovascular dysfunction is in neurodegenerative disease. One popular view is that cerebrovascular dysfunction is a consequence of a primary dysfunction, such as dysfunction of the neuronal microchondria. In this case, the inability of neurons to manage their energy needs results in their demise [32], which in turn activates glial cells. In contradistinction, other studies point to the causal effects of cerebrovascular dysfunction (resulting from inflammation) on neurodegenerative disease, with neuronal death occurring as a result of vessel constriction or stiffening, or glial dysfunction [11].

Whatever the true series of events that precede neuronal death in the prodromal phase of neurodegenerative diseases, the astrocytes are in an activated state. It is believed that this activation compromises their ability to maintain neurovascular coupling, creating a potential biomarker for detecting early-stage neurodegenerative diseases, for example using generative embedding.

Reasons why astrocytes may be unable to maintain neurovascular coupling when they are in an activated state include:

- Astrocytic gap junction protein Connexin 43, which mediates inter-astrocyte communication, is inhibited during inflammation, reducing the efficacy of

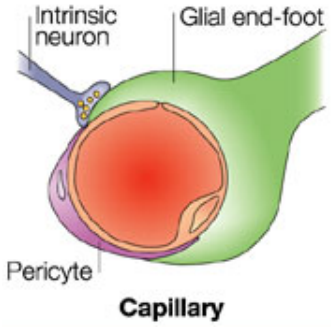

Figure 5.3: Functional unit for capillaries: neurons and glial cells have processes on the outer surface of blood vessels, allowing them to act together with vascular cells to control CBF. Figure reproduced from Iadecola, 2004, Box 2 [82].

astrocytes to co-operate when producing vascular changes [98].

- COX2, an enzyme involved in generating vasoactive molecules in response to neuron-to-astrocyte signalling, is upregulated by pro-inflammatory cytokines, resulting in imbalanced neurovascular mechanisms [69].

5.1.2 Predictions

Using the hypothesised series of events leading to neurodegenerative disease onset and progression described in the previous subsection, and in particular the consequences of these events for neurovascular coupling/uncoupling, it is possible to make predictions about the disease-state discriminabilities that might be observed using the traditional analysis approach (in this chapter) and the generative embedding approach (in the next chapter). For example, because the multimodal couplings of amplitude and power shape features are all positive, only differences in the dependences of the positive couplings provide discriminatory power in traditional analyses. For this reason, the discriminability of early-stage neurodegenerative diseased and healthy neuroimaging signals should theoretically be greater than the discriminability of established neurodegenerative-diseased and healthy signals when they are analysed using the traditional method. In contradistinction, the discriminabilities of neurodegenerative-diseased and healthy neuroimaging feature spaces should theoretically be greater during the established

neurodegenerative disease stage than at any earlier stage when they are analysed using generative embedding. This is because generative embedding is likely to generate multimodal spaces with negative couplings, and the most discriminable feature spaces are those with negative couplings between neuronal and hemodynamic features that are both modulated by disease, as shown in Fig.5.4.

When one is instead concerned with the **added** discriminability that imaging the neuronal activity and hemodynamics concurrently has over imaging them separately, the value should theoretically be greatest for earlier stages of disease, regardless of whether negative couplings are captured or not. This is because it is neurovascular uncoupling rather than coupling strength that determines whether there is a synergistic effect, as explained in Fig.5.5.

The only stage of disease progression that may benefit from concurrent imaging of the neuronal activity and hemodynamics more than the early stage is the very late stage, when impairment of neuronal activity exceeds that for hemodynamics because a cycle of neuronal deterioration has developed. However, as previously mentioned, the focus of this work should be to diagnose and treat early-stage neurodegenerative disease as the value of delaying cognitive decline is much greater than slowing it down once it is already established.

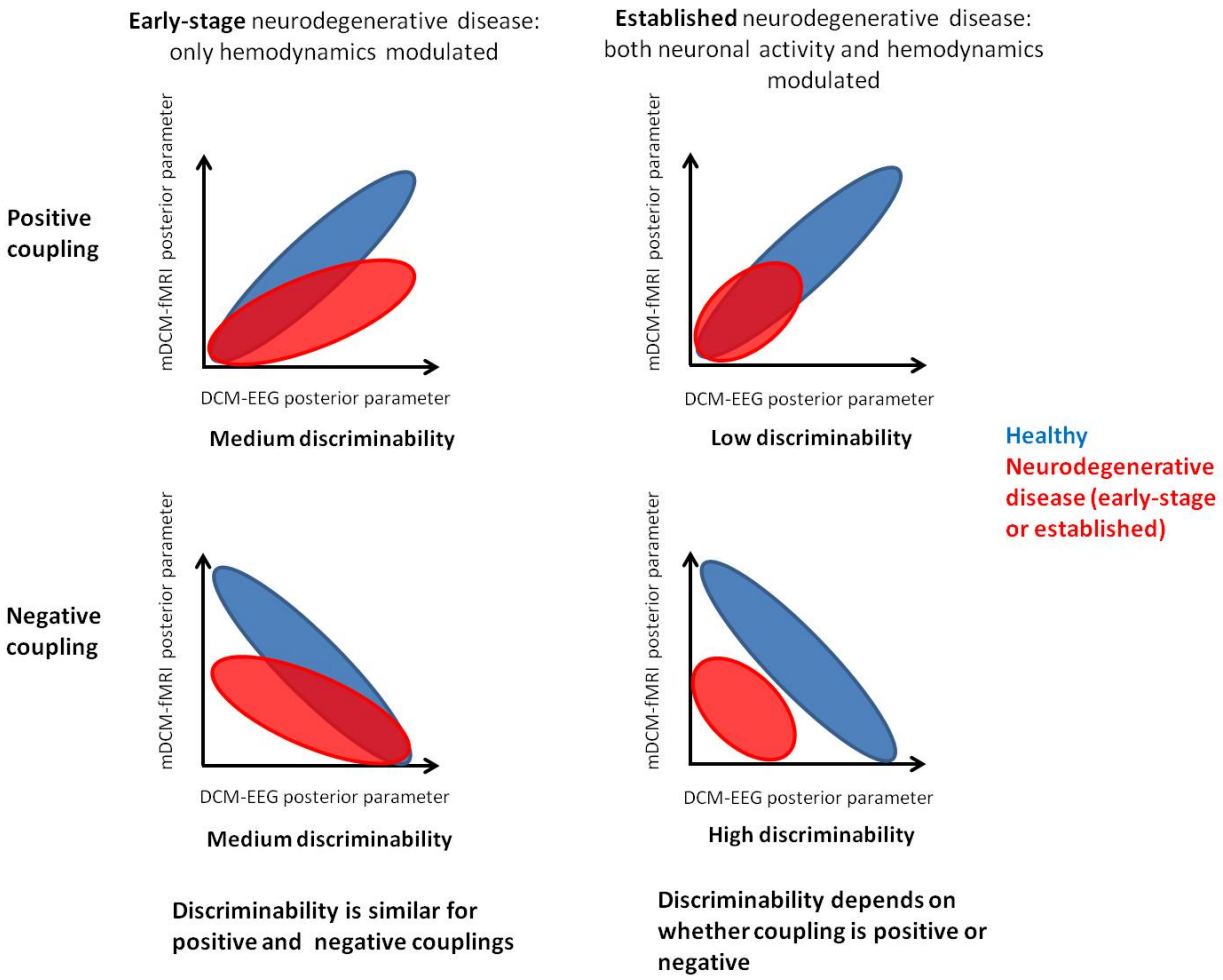

Figure 5.4: Explanation of the hypothesis that the statistical power with which neurodegenerative-diseased and healthy neuroimaging signals can be discriminated should theoretically increase with disease progression when the feature space has negative multimodal coupling, but decrease otherwise. This hypothesis is explained here using pairs of multimodal DCM posterior parameters as this chapter is preparatory for Chapter 6, which aims to demonstrate the utility of generative embedding for clinical and drug development applications.

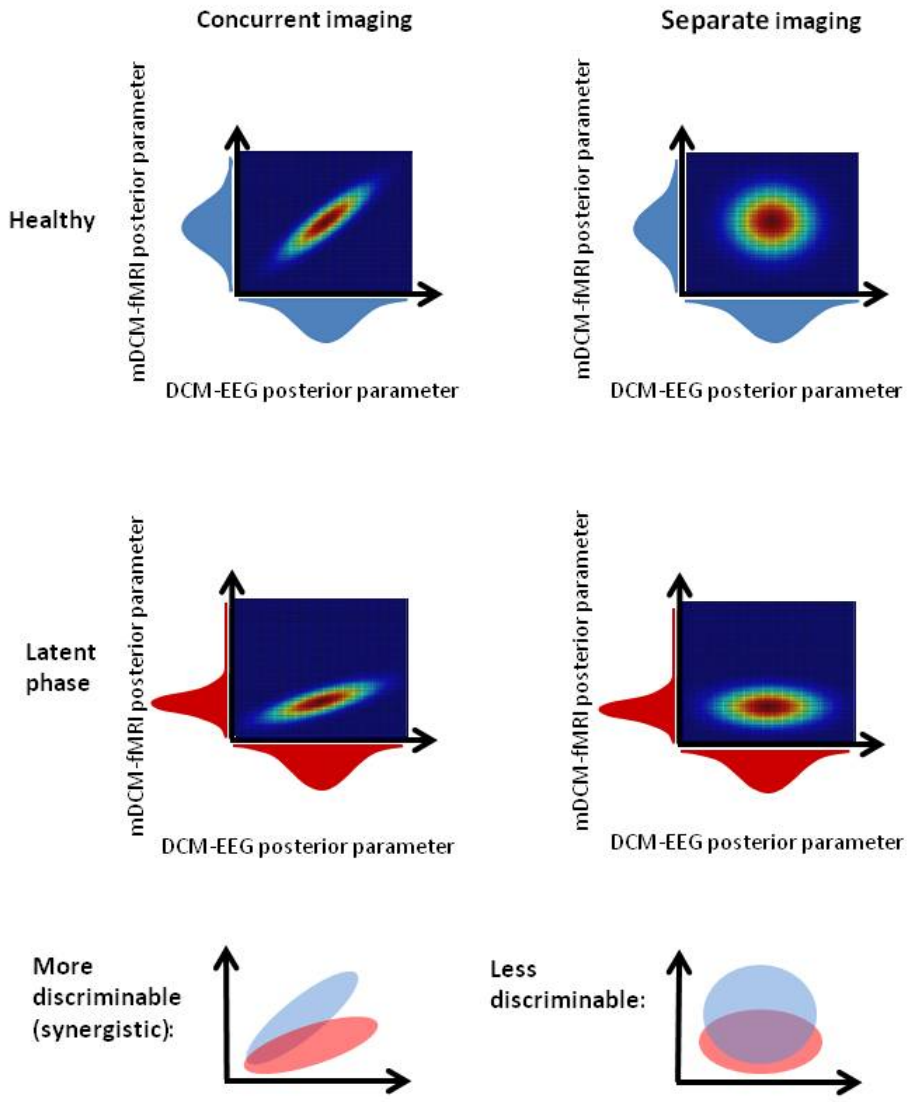

Figure 5.5: Explanation of the hypothesis that if there are changes in the couplings between the DCM-EEG and mDCM-fMRI posterior spaces between two conditions (e.g. healthy and latent neurodegenerative diseased brain) then it is possible to capitalise on these couplings to improve the discriminability of the conditions. Shown here in 2D only, i.e. between one DCM-EEG parameter and one mDCM-fMRI parameter rather than between their multivariate spaces, it is possible to see that when the EEG and fMRI signals are acquired concurrently, their joint posterior parameter probability densities for healthy and latent disease phase patients are more discriminable (bottom left) than when they are acquired separately and the differences in the multimodal couplings are not available. This is because when the signals are acquired separately, the trial-wise intra-subject variations in the couplings are not captured (bottom right). Notice that the magnitude of the (theoretical) EEG parameter is the same for both the healthy and diseased conditions here, in keeping with the hypothesis that cerebrovascular dysfunction occurs before neuronal impairment.

5.2 Methods

5.2.1 Experimental data collection

Lipopolysaccharide (LPS), a lipid of negative-gram bacteria that elicits a strong immune response[1], was chosen as the pharmacological agent for inducing early neurodegenerative-disease like glial activation in this DPhil. Although this immune response is systemic rather than central, there is evidence that a systemic inflammatory response induces a central response [150]. This is because systemic inflammation results in rapid cytokine (inflammatory messenger protein) production. These proinflammatory cytokines, which include interleukin-1, (IL-1), interleukin-6 (IL-6), and tumour necrosis factor-α (TNF), circulate in the blood and communicate with the microglia (macrophages located in the brain), which themselves then become activated, resulting in a central inflammatory response. The evolutionary reason for this phenomenon is to induce sickness behaviours that aid recovery, such as rest [150].

5.2.1.1 Experimental protocol

Male Sprague-Dawley rats (n = 4 per group, 230-370 g; Harlan, UK) were anesthetised with 2-3% Isoflurane in a mixture of nitrous oxide/oxygen (70%/30%). They were then given an i.p. injection of LPS (5mg/kg delivered in 1ml/kg solution for disease animals) or the equal volume of saline (for healthy

[1]The acute immune response elicited by LPS is classical activation, which is one of four immune responses. Two of the immune responses are proinflammatory, "killer" states: innate (induced by pattern recognition receptors when they encounter a pathogen), and classical (induced naturally by IFN-γ after innate activation, or by external stimuli such as LPS). The other two immune responses are anti-inflammatory, repair and immunosuppression states that occur afterwards: alternative activation (induced by IL-4 and IL-3), and acquired deactivation (induced by TGF-β, IL-10 and apoptic cells). Chronic inflammation is believed to occur when classical and alternative activation co-exist [34]. For example, in late stage Alzheimer's disease, classical and alternative activation are stimulated to co-exist by A-β: the A-β plaques stimulate classical activation, causing glial cells (including astrocytes) to prioritise degradation of A-β over their other functions, such as neuronal protection. The resulting neuronal damage results in alternative activation, as glial cells make reparative attempts. It is unknown how plastic glial cells are, and therefore whether individual cells cycle their activation state between classical and alternative, or whether different cells can manifest just one of the different activation states [34].

animals). Subsequently, Isoflurane was reduced to 2% and the animals were tracheotomised and artificially ventilated. The left femoral artery was cannulated for withdrawal of blood samples in order to monitor blood pressure and blood gas levels. The animals were then placed in a stereotaxic frame, a midline incision was made in the scalp and the area above barrel cortex was thinned to translucency. A burr hole was drilled through to the barrel cortex (3.5 mm posterior and 2.5 mm lateral to Bregma) and the recording electrode was inserted approximately 3mm into the barrel cortex via this hole. A Speckle Contrast Imager (moorFLPI, Moor Instruments) was then focused on the translucent area of skull and the stimulating electrodes were placed in the whisker pad subcutaneously. Once three hours had passed since the LPS/saline injection (to ensure peak effect of LPS on the brain [160]), the contralateral whisker pad was stimulated electrically using a Grass stimulator (Grass Instrument Co., Quincy, Mass.) triggered to deliver 60 trains of 0.3ms electrical pulses at a 10Hz repetition rate lasting 2s each. The resulting LFPs[1] measured by the recording electrode were sampled at 10,000Hz by Spike2 software (Cambridge Electronic Design, UK) after amplification by CED 1902 and 1401 and noise-reduction by Humbug (Digitimer Ltd.). The resulting CBF responses were concurrently measured by the Speckle Contrast Imager at 25Hz and processed using moorFLPI Review software. Throughout data collection, 1-1.5% Isoflurane was used as maintenance anaesthesia. Core body temperature was regulated and maintained at 37°C throughout the experiment using a heating mat. All experiments were approved by the UK Home Office Animals (Scientific Procedures) Act (1986) in line with EU directives on animal experimentation.

[1]LFPs rather than EEG were recorded in this dataset because the existing set-up of the lab supported this more easily. LFPs and EEG are different measures of the same neuronal activity, and therefore any insights gained from the LFPs are translatable to EEG recordings too.

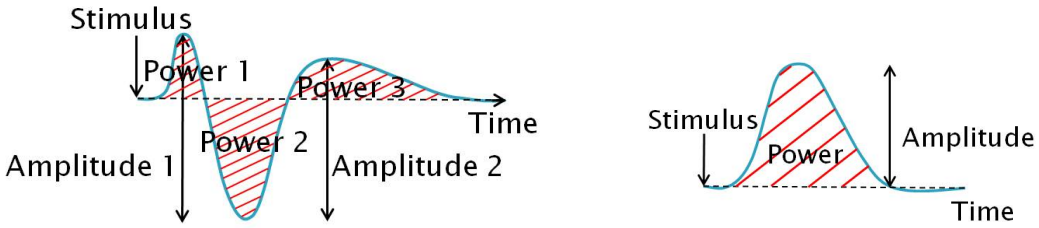

Figure 5.6: Diagrams of the shape features extracted from each LFP (left) and CBF (right) response to the stimuli. The CBF response shown here is an expected response to a train of stimuli, whereas the LFP response is the expected response to a single stimulus. To obtain a single value for each LFP shape feature for a train of stimuli, the values found for the constituent single responses were summed.

5.2.2 Simple analysis: linear regression of shapes features

To gain an idea of the ability of Dartmouth Brain Imaging Group's system for detecting brain pathologies from the aspects of neurovascular coupling inferrable using the traditional analysis (as their system performs linear classification of multimodal shape features), the traditional analysis was used to identify whether significant differences exist in the linear relationships of the multimodal shape features. The shape features extracted in this work are the maximum amplitude and power of the trial-wise LFP and CBF responses to the stimuli, as shown in Fig.5.6.

Using R's "lme" function[1], the following mixed-effects regression model was fitted to the CBF shape features, $\boldsymbol{y}$, recorded under each disease state separately:

$$\boldsymbol{y} = \boldsymbol{A}(\boldsymbol{\beta} + \boldsymbol{b})\boldsymbol{x} + \boldsymbol{\epsilon}, \tag{5.1}$$

where $\boldsymbol{b} \sim N(0, \boldsymbol{\Sigma})$ and $\boldsymbol{\epsilon} \sim N(0, \boldsymbol{\sigma}^2)$. In this model, $\boldsymbol{\beta}$ are the fixed-effect coefficients (modelling the linear couplings between the multimodal shape features of the population of rats under that disease state), $\boldsymbol{b}$ are the normally distributed rat-specific random-effects (zero-mean rat-specific changes to the population linear coupling under that disease state) and $\boldsymbol{x}$ are the rat-specific regressors (all the

[1] "lme" fits regression models by maximising the marginal likelihood of the model (integrated over the random effects) with the constraint that the random effects are uncorrelated.

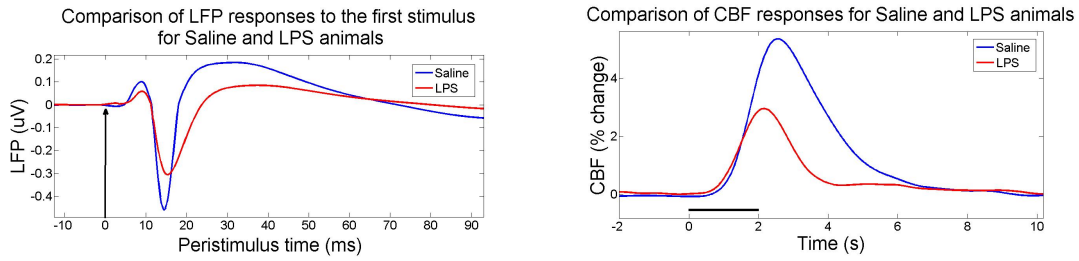

(a) (b)

Figure 5.7: Comparison of the average (a) LFP responses to the first stimulus in the stimulation train, and (b) CBF responses to the train of stimuli for the LPS (diseased, red) and control (healthy, blue) animals. Notice that the average LFP for the LPS animals is not only smaller but also slightly delayed, whereas the average CBF response for the LPS animals is not only smaller but also shorter in duration. It is not possible to comment on the relative timings of the CBF responses because of errors with the calibration of the Speckle Contrast Imager which meant that the stimulation times had to be guessed from the responses.

corresponding LFP shape features). Because no intercept coefficient is included in this model, the condition that when no neuronal activity is detected then no CBF response is observed, is imposed.

P-values describing the significances of the differences of the $\boldsymbol{\beta}$ estimated for the LPS data compared to the control data were calculated using the contrast [1 -1].

5.3 Results

The average LFP and CBF responses are shown in Fig.5.7. It is clear that the diseased animals have smaller, slower LFPs and smaller, shorter CBF responses compared to the control animals. Unfortunately, because of calibration errors with the Speckle Contrast Imager, it is not possible to know from these data if the hemodynamics of the diseased animals were also slower.

The regression model fits for each pair of multimodal shape features are shown in Fig.5.8 along with the individual datapoints, clustered in colour by rat. There are two clear observations to discuss from these plots:

- The amplitudes and powers of both the neuronal activity and hemodynamics

shape features are considerably smaller for the LPS group than the control group (the lengths of the lines reflect the range of values extracted from all the individual responses). This indicates that the LPS induced neuromodulatory effects similar to those hypothesised to occur in neurodegenerative disease (neuronal impairment/death caused by hemodynamic impairment), but not the effects specific to early-stage neurodegenerative disease, i.e. modulated hemodynamics only.

- Only the couplings of LFP power 2 and 3 with CBF power (Fig.5.8(i) and (j)) were found to be significantly different at the 95% significance level, suggesting that these are the shape features that ought to carry large weights in the linear classifier incorporated in Dartmouth Brain Imaging Group's prototype product for detecting established neurodegenerative disease. These couplings are significantly different between the disease states because the hemodynamics are modulated by the LPS to a greater extent than the neuronal activity is.

5.4 Discussion

It is clear from Fig.5.8 that there was considerable inter-subject and intra-subject variation in all of the shape features extracted from the LFP and CBF signals, and these sources of variation will have weakened any statistical power for detecting significant group differences between the multimodal shape features. Both of these sources of variation may be attributable to experimental technique:

- Despite both the anaesthesia and LPS neuromodulations having inconstant biological effects, it was necessary to analyse all of the data together rather than separate them into different timepoints. This was because the set-up of the experimental laboratory made it very difficult to collect LFP data without considerable noise, meaning that larger sample sizes were needed for

sufficient statistical power to be achieved. Analysing all of the data together in this way increased the intra-subject variations in the linear relationships between the multimodal shape features and, in doing so, reduced the statistical power of any inter-group differences.

- The large inter-subject variations may be an artifact of my (relatively inexperienced) experimental technique rather than a true reflection of biological inter-subject variation.

The consequence of these extraneous sources of variation may be that the linear relationships between other multimodal shape feature pairs, especially LFP power 3 vs. CBF amplitude (Fig.5.8(e)), are in fact significantly different for the healthy and diseased animals and therefore ought to carry large weights in the linear classifier too.

In Chapter 6, DCMs will be fitted to a dataset that does not exhibit the high levels of inter-subject and intra-subject variation or errors in the timing information of the hemodynamics data that this preliminary dataset does. The resulting DCM posterior parameter estimates will then be embedded into information-theoretic statistics in order to quantify their discriminabilities. As previously mentioned, this generative embedding approach should provide spaces with negative multimodal couplings, increasing the discriminability of the disease states beyond that achieved in this chapter using the traditional analysis. Therefore, although feature spaces including LFP power 2 and 3 and CBF power were found to be significantly different in this chapter, the conclusion made from the generative embedding analysis carried out in the next chapter is expected to be that spaces capturing aspects of neurovascular coupling other than amplitude and power ought to be used when classification or regression is carried out on generative spaces instead.

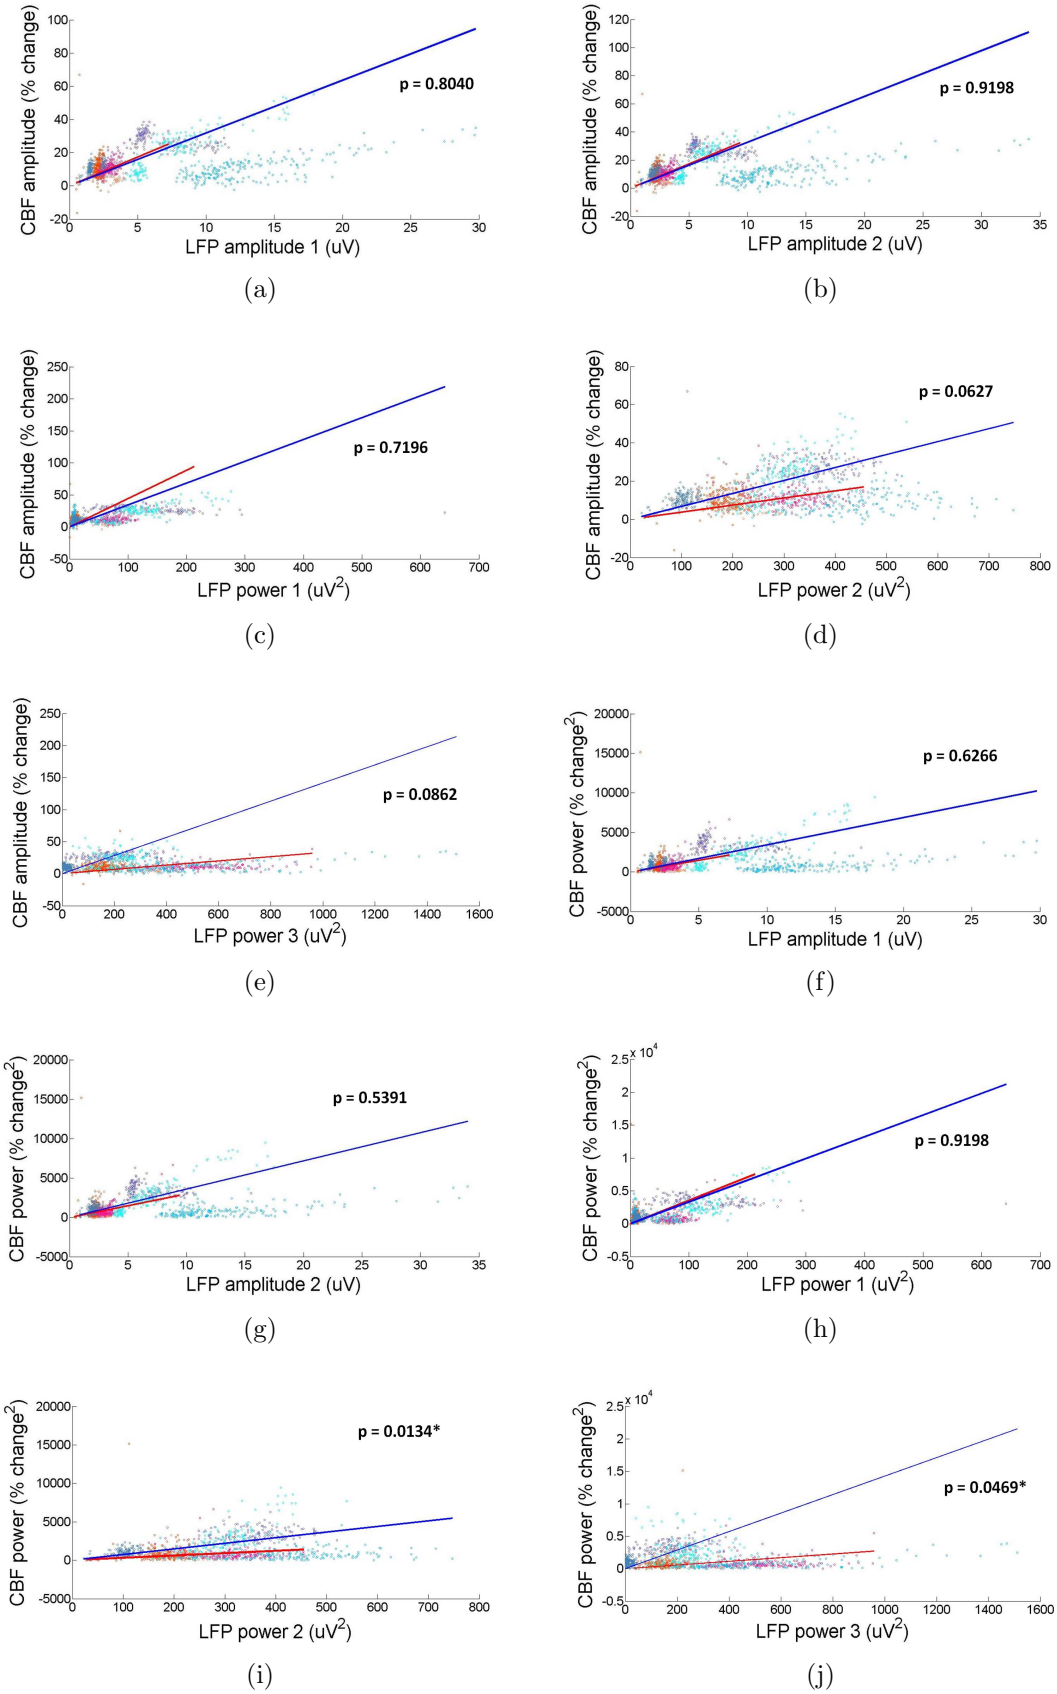

Figure 5.8: Mixed effects regression lines between the multimodal shape features for disease (red) and healthy (blue) animals. The individual datapoints are also shown, clustered by colour for each rat (the shades of red correspond to disease rats and the shades of blue correspond to healthy rats). The couplings were found to be significantly different only for the bottom two pairs of multimodal shape features (LFP power 2 and 3 with CBF power).

Chapter 6

Can we exploit neuromodulatory effects to better diagnose neurodegenerative disease?

In Chapter 5 a preliminary study was carried out in order to investigate whether significant differences could be found between healthy and early-stage neurodegenerative-diseased neuroimaging signals using the traditional analysis approach (linear regression of shape features). The experimental protocol did not induce the intended neuromodulatory effects as neuronal activity was modulated by the LPS (Lipopolysaccharide, used to elicit a strong immune response) as well as the hemodynamics - a profile of the functional unit that is more similar to that of established neurodegenerative disease, where some cognitive decline has occurred, than to early-stage neurodegenerative disease. Nevertheless, significant differences were found in the linear couplings of two of the LFP power features with CBF power in these preliminary data, suggesting that they ought to strongly inform the linear classifier used in Dartmouth Brain Imaging Group's prototype product for detecting established neurodegenerative diseases.

The aim of this chapter is to investigate the potential of generative embedding

for diagnosing, monitoring and understanding neurodegerative disease, by:

1. Quantifying the discriminability of early-stage neurodegenerative-diseased and healthy neuroimaging signals, as the more discriminable the disease states are, the more accurate their classification (for disease diagnosis) and the more significant their regression statistics (for disease monitoring and understanding) are likely to be.
2. Investigating whether the discriminability of early-stage neurodegenerative-diseased and healthy neuroimaging signals is increased when neuronal activity and hemodynamics are imaged concurrently rather than separately, which would support the development of hybrid neuroimaging systems for clinical and drug development applications.

This potential is itself explored using generative embedding.

6.1 Description of neurodegenerative disease dataset

The data analysed in this chapter were collected by Rosengarten et al. [164]. These data were chosen because Rosengarten et al. reported significant differences in the linear relationships of shape features extracted from the neuroimaging timeseries of healthy and LPS animals similar to those found in the preliminary dataset. However Rosengarten et al.'s dataset was collected using a correctly calibrated system (unlike the preliminary dataset, which suffered timing errors) and it has the advantages of larger sample size and also much less noisy data (meaning that separate analyses can be carried out on the data according to the time since the saline/LPS injection was administered), increasing the statistical power of any effects detected.

Rosengarten et al.'s motivation for comparing neurovascular coupling in control

and LPS rats was to investigate the effect of sepsis (whole-body inflammatory state to severe infection) on neurovascular dysfunction, as survivors of severe septic illness are often reported to suffer cognitive disability [93]. This dataset contains concurrent EEG and CBF responses to forepaw stimulation in three groups of Alpha-Chloralose anaesthetised rats: ten animals given saline i.v. (control animals), ten animals given 1mg/kg LPS i.v., and ten animals given 5mg/kg LPS i.v. The EEG and CBF responses were recorded from somatosensory cortex at six timepoints following the i.v. injection: 30 minutes, 1 hour, 2 hours, 3 hours, 4 hours and 4.5 hours. The stimulation paradigm comprised ten trains of 2Hz stimuli lasting 30 seconds and separated by 30 seconds of rest.

Similar to the work carried out in Chapter 5 on the preliminary dataset, Rosengarten et al. used simple linear regression of average shape features (N1-P1 SEP amplitudes and CBF percentage change) to compare neurovascular coupling in the different conditions. Figure 6.1 compares the average amplitude features of the three groups over the recording period. Both the SEP and CBF amplitudes reduced to stable levels within 1 hour after 1mg/kg LPS injection, meaning that their linear coupling was maintained at all timepoints under this dosage of LPS. However, for the 5mg/kg LPS group, the CBF amplitudes dropped to stable levels 1 hour after injection, whereas SEP amplitudes took 3 hours to stabilise, meaning that the hemodynamics were modulated to a greater extent than the neuronal activity between these timepoints. This neurovascular uncoupling for the 5mg/kg group between 1 and 3 hours is similar to the significant uncoupling observed in the preliminary dataset, where the hemodynamics amplitude/power shape features were reduced by 5mg/kg LPS to a greater extent than the neuronal activity shape features (some significantly so), although in the preliminary dataset this uncoupling was observed to (still) occur more than 3 hours after injection.

As neurovascular uncoupling was only observed to occur in Rosengarten et al.'s data for the 5mg/kg LPS rats between 1 and 3 hours after LPS injection, only

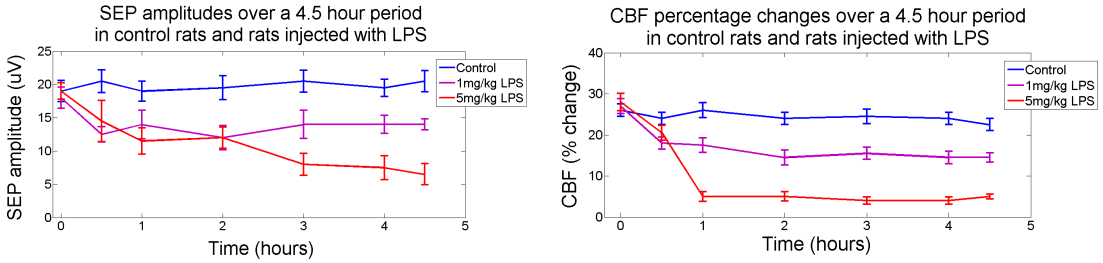

Figure 6.1: Comparison of (a) SEP and (b) CBF amplitudes of the three conditions over time [164]. SEP amplitude of the two LPS groups is significantly different from the control groups at all timepoints after injection, however the SEP amplitude of the 5mg/kg LPS group is only significantly lower than the 1mg/kg group 3 hours after injection or longer. CBF amplitude is significantly different between all three conditions from 1 hour after injection onwards.

these data were requested[1]. Unfortunately, the data analysed in Rosengarten et al., 2006, has since been destroyed, and so the group provided me with another dataset that they had collected using the same experimental protocol. These data were used to meet the aims of this chapter.

6.2 Generative Embedding: Information-theoretic statistics

The performance of any algorithm used to analyse neuroimaging signals in order to diagnose, monitor and understand early-stage neurodegenerative disease will benefit from greater discriminability of the disease-state imaging signals. For this reason, the generative embedding algorithm used in this chapter to explore the potential of generative embedding for these purposes embeds DCM posterior parameter estimates into information-theoretic statistics, as these are sensitive to all orders of correlation, i.e. discriminability in its truest sense. Specifically, to meet Aim 1, which is to quantify the discriminability of neurodegenerative disease-state neuroimaging signals, the mutual information (MI) of multimodal

[1]It was concluded that it is not possible to induce the intended neuromodulatory effects of early-stage neurodegenerative disease using LPS, as discussed in the Conclusions Chapter.

DCM posterior parameters spaces are calculated. To meet Aim 2, which is to investigate the value of concurrently imaging neuronal activity and hemodynamics, their synergy is calculated.

6.2.1 Information Theory

Information theory is a branch of probability theory first developed by Shannon in the 1940s to quantify limits on the information in signals of communication [173]. Information is based on the notion of the entropy of a random process X:

$$\mathrm{H}(X) = \mathrm{E}_{\mu_x} \log_2 \frac{1}{\mu_x(x)} = -\int_x \mu_x(x) \log_2 \mu_x(x) \mathrm{d}x, \tag{6.1}$$

where $\mathrm{H}(X)$ is the entropy of random variable X with density $\mu_x(x)$ and E is its expectation. Entropy measures the average uncertainty in a random variable, which, owing to the use of logarithms to base 2 (so that it is measured in bits), is equal to the number of "yes/no" questions necessary to describe it, so that 1 bit of information decreases uncertainty of the random variable by half (or increases discriminability of the random variable by two).

MI is a measure of the dependence between two sets of random variables **A** and **B**:

$$\mathrm{I}(\boldsymbol{A};\boldsymbol{B}) = \mathrm{H}(\boldsymbol{A}) - \mathrm{H}(\boldsymbol{A}|\boldsymbol{B}) = \int_a \int_b \mu(\mathrm{a,b}) \log \frac{\mu(\mathrm{a,b})}{\mu_a(a)\mu_b(b)}, \tag{6.2}$$

where μ(a,b) is the joint density of **A** and **B**.

Because the entire joint density is used to calculate MI, it is sensitive to all orders of correlation, i.e. to nonlinear relationships, in contrast to linear regression, which only considers second-order dependences (covariances) of density averages.

MI is symmetric in **A** and **B**, can never be negative and is only zero if **A** and **B** are independent (so that $\mathrm{H}(\mathbf{A}|\mathbf{B}) = \mathrm{H}(\mathbf{A})$). These are consequences of the fact that MI is a special case of the more general relative entropy, which is the

Kullback–Leibler distance between two probability mass functions p and q [37]:

$$\mathrm{D}(p||q) = \mathrm{E}_p log \frac{p(\mathrm{a})}{q(\mathrm{b})}, \tag{6.3}$$

which is non-negative and only zero when $p(\mathrm{a}) = q(\mathrm{a})$.

6.2.1.1 Synergy

The neuroscience community extended the application of MI to understand how neurons encode information: do they convey redundant information, independent information, or is the information they convey synergistic, i.e. do correlations in their firings result in more information being carried than the sum of their individual informations [66][169][134][143]? These questions can be answered by calculating the MI to test for independence, and the synergy to test for the redundancy or synergy of the firings of neighbouring neurons:

$$\mathrm{Synergy}(\boldsymbol{S};\mathrm{A,B}) = \mathrm{I}(\boldsymbol{S};\mathrm{A,B}) - \mathrm{I}(\boldsymbol{S};\mathrm{A}) - \mathrm{I}(\boldsymbol{S};\mathrm{B}), \tag{6.4}$$

where I($\mathbf{S}$;A,B) is the MI between the stimulus set $\mathbf{S}$ and the joint distribution of the firings of neurons A and B (i.e. the probability that $\mathbf{S}$ can be accurately discriminated from A and B), and I($\mathbf{S}$;A) and I($\mathbf{S}$;B) are the MI between the stimuli and the firings of neurons A and B, respectively. A negative synergy indicates that there is redundancy in the information carried by neurons A and B, whereas a positive synergy indicates that they convey information synergistically.

Panzeri et al. proposed that the application of synergy be extended to quantify the value of EEG-fMRI [144]. Ostwald et al. carried out such an investigation, finding that combining EEG and fMRI time-domain features improved the significance of the high and low checkerboard stimulation contrast [139].

6.2.2 Bias in information-theoretic calculations

In order to calculate the MI or synergy between **A** and **B** for explaining **S**, their probability densities must first be estimated. Unfortunately, sampling fluctuations in (small) datasets result in both systematic (bias) and statistical (variance) errors in the density estimates, known as finite sample effects, as shown in Fig.6.2.

All entropy estimates are negatively biased, with the bias being about **S** times larger for conditional estimates, i.e. those conditioned on **S**, such as H(**A**|**S**), than non-conditional estimates such as H(**A**). This is because the conditional estimates are made from **S** times less data. The overall effect of these negative entropy biases are positive biases in the MI values and therefore a positive bias in the overall synergy value (because the positive bias in I(**S**;**A**,**B**) will be larger than the positive bias from the sum of I(**S**;**A**) and I(**S**;**B**)).

The bias and the standard deviation decrease when more data is available to estimate the densities from because the fluctuations in the estimated probabilities become smaller. Nevertheless, the most commonly used method for estimating the probability density of a random variable, namely binning its observations, exacerbates the effect of there appearing to be more structure in the data than there actually is, increasing the bias. There are two options for improving density estimates when the amount of data remains low:

1. Apply a bias correction scheme to the binning strategy.
2. Use a binless strategy for estimating the densities.

6.2.2.1 Choosing an algorithm to estimate densities with low bias

Papana, 2009, compared the performances of the following commonly used binning and binless estimators of MI [147]:

1. Binning: histograms of fixed, equidistant bin size (the most common density estimation method, as mentioned in the previous section)

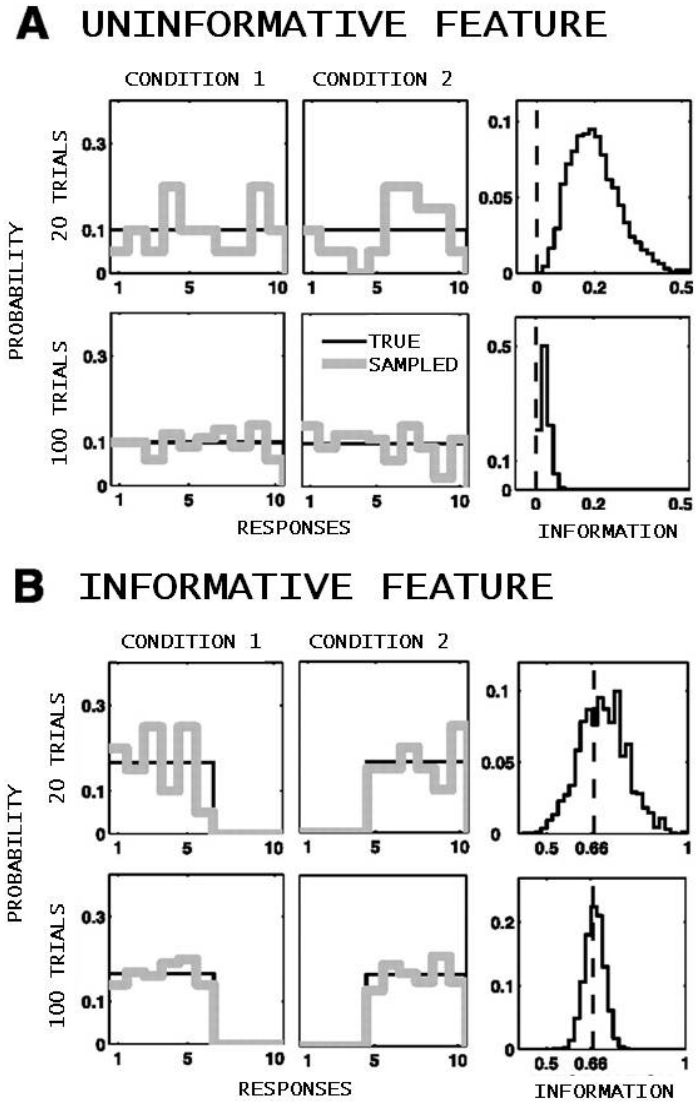

Figure 6.2: A: Response probability histograms from simulations of an uninformative feature, whose response on each trial follows the uniform distribution ranging from 1 to 10 for both conditions (no discriminability between conditions), sampled from 20 and 100 trials per condition. The black horizontal line is the true uniform response distribution. Right: distribution (over 5,000 simulations) of the information values obtained with 20 (top) and 100 (bottom) trials per stimulus, respectively. As the number of trials increases, both the bias and standard deviation decrease. The dashed vertical line in the right columns indicates the true value of the information carried by the simulated features (here 0, as they are not discriminable). B: Response probability histograms from simulations of an informative feature, whose response on each trial has uniform probability for being 1 to 6 in condition 1, and 5 to 10 in condition 2 (conditions are discriminable, hence positive MI). Right: once again, as the number of trials increases, both the positive bias and standard deviation decrease. Figure reproduced from Panzeri et al., 2007, Figure 1 [145].

2. Binning: histograms of adaptive bin size so that bin occupancy is approximately uniform (see Appendix D for a literature review of the bias correction schemes used to achieve this transformation or correct for the bias in the binning method in some other way)
3. Binless: k-nearest neighbours, which uses the distances of the k^{th}-nearest neighbour to each datapoint in the joint density to estimate the joint and marginal densities about that point
4. Binless: kernels, which are centred at datapoints to obtain weighted distances in the joint and marginal spaces

When using any of these estimators there is at least one free parameter that must be set: when using histograms it is the number of bins, when using k-nearest neighbours it is k, i.e. how many neighbouring points are used to estimate the local density, and when using kernels it is the bandwidths of the kernels. The best values for these parameters depend on the complexity, noise level and size of the data.

Because the analytical MI of certain bivariate distributions are known, the accuracy and stability of these different estimators can be tested. An example of such a distribution is the bivariate normal distribution:

$$\mathrm{MI} = -0.5 \log (1 - \rho^2), \tag{6.5}$$

where ρ is its correlation coefficient.

Using Monte Carlo simulations of bivariate normal distributions, Papana, 2009, found that k-nearest neighbours gave the most stable and accurate density estimates whilst also being least affected by its free parameter, k, i.e. it is least vulnerable to calculating a poor entropy estimate owing to the choice of k. For these reasons, k-nearest neighbours was chosen as the estimator for Rosengarten et al.'s LPS dataset, because the complexities of the DCM posterior parameter spaces are unknown and may be very variable, meaning that different k values might be

optimal for different parameters.

6.2.2.2 k-nearest neighbours

k-nearest neighbours was developed by Kraskov et al., 2004, to calculate the MI between two random variables A and B [91]. It is based on the same density estimation as kernels:

$$\mathrm{p(a)} \approx \frac{\mathrm{k}}{\mathrm{N}V}, \tag{6.6}$$

where p(a) is probability density at datapoint a, V is the volume surrounding a, k is the number of datapoints inside V and N is the total number of datapoints.

The difference between density estimation using kernels and k-nearest neighbours is that in kernel density estimation V is set (by the kernel bandwidth) and the corresponding k are calculated, whereas in k-nearest neighbours k is set and the corresponding V are calculated[1].

k-nearest neighbours is an extension of the Kozachenko-Leonenko estimate for univariate entropies and exploits the fact that the entropy of random variable A is the negative average of log μ(a) (see Equ.6.1) to calculate an unbiased entropy estimate from an unbiased estimate of log $\mu(\mathrm{a}_i)$ [90]:

$$\hat{\mathrm{H}}(\mathrm{A}) = -\psi(\mathrm{k}) + \psi(\mathrm{N}) + \frac{1}{\mathrm{N}} \sum_{i=1}^{\mathrm{N}} \log \epsilon(i), \tag{6.10}$$

[1]This k-nearest neighbours algorithm for density estimation uses the same concepts as the k-nearest neighbour algorithm used in classification (sometimes known as Bayes classifier): to classify unknown sample a to one of the **w** possible classes,

$$\text{Likelihood, } \mathrm{p(a}|w_i) = \frac{\mathrm{k}_i}{\mathrm{N}_i \mathrm{V}}, \tag{6.7}$$

$$\text{Prior, } \mathrm{P}(w_i) = \frac{\mathrm{N}_i}{\mathrm{N}}, \tag{6.8}$$

$$\text{Posterior, } \mathrm{P}(w_i|\mathrm{a}) = \frac{\mathrm{p(a}|w_i)\mathrm{P}(w_i)}{\mathrm{p(a)}} = \frac{\mathrm{k}_i}{\mathrm{k}}. \tag{6.9}$$

Therefore, a is assigned to the class most greatly represented within V.

where ψ is the digamma function and ϵ(i) is the length that extends from the ith datapoint to its kth-nearest neighbour datapoint.

Equ.6.10 can be derived by considering, for each datapoint a_i:

- p_i to be the mass of the length ϵ(i)
- $P_i(\epsilon(i))d\epsilon$ to be the probability that the kth-neighbour lies within [$\epsilon(i)$, $\epsilon(i) + d\epsilon$] from a_i (and that k-1 other points lie at shorter distances, and N-k-1 at larger distances):

$$P_i(\epsilon(i))d\epsilon = k\left(\frac{N-1}{k}\right)\frac{dp_i}{d\epsilon}d\epsilon \cdot p_i^{k-1} \cdot (1-p_i)^{N-k-1} \tag{6.11}$$

The expectation of log p_i(a) can then be written as:

$$\begin{aligned}\int_0^\infty P_i \log p_i \, d\epsilon = k\left(\frac{N-1}{k}\right)\int_0^1 p(k-1)(1-p)^{N-k-1}\log p \, dp, \\ = \psi(k) - \psi(N).\end{aligned} \tag{6.12}$$

Then, by assuming that $\mu(a_i)$ is constant across the length, the mass p_i can be approximated by $\epsilon(i)\mu(a_i)$ (volume x density), giving:

$$\log \mu(a_i) \approx \psi(k) - \psi(N) - \log \epsilon(i), \tag{6.13}$$

resulting in the Kozachenko-Leonenko entropy estimate (Equ.6.10).

Extending this method of entropy estimation to an m-dimensional space gives:

$$\widehat{H}(A_1, A_2, ..., A_m) = -\psi(k) + \psi(N) + \frac{m}{N}\sum_{i=1}^{N} log\, \epsilon(i), \tag{6.14}$$

where the $\boldsymbol{\epsilon}$ are now the side lengths of hyper-cubes assumed to have constant density (see Fig.6.3).

In Kraskov et al., the MI between two random variables (A,B) is calculated

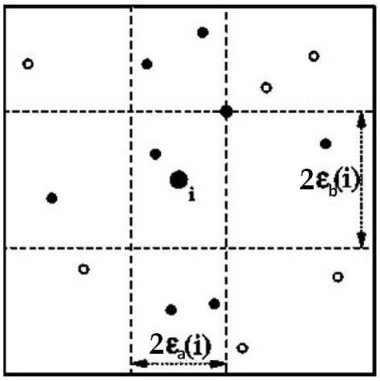

Figure 6.3: Example 2D density showing the square about datapoint i in which the mass is modelled as constant in k-nearest neighbours density estimation. $\epsilon(i)$ is set to two-times the maximum of the marginal distances separating i from its k-nearest neighbour, i.e. ϵ(i) = max($2\epsilon_{\mathrm{a}}$(i), $2\epsilon_{\mathrm{b}}$(i)). Figure reproduced from Kraskov et al., 2004, Figure 1a [91].

using k-nearest neighbour estimates for H(A) + H(B) - H(A,B) [1] [184]:

$$\widehat{\mathrm{I}}(\mathrm{A,B}) = \psi(\mathrm{k}) - \frac{1}{\mathrm{N}}\sum_{i=1}^{\mathrm{N}}(\psi(\mathrm{n}_a(i)) + \psi(\mathrm{n}_b(i))) + \psi(\mathrm{N}), \tag{6.16}$$

where $\mathrm{n_a}$ and $\mathrm{n}_{\boldsymbol{b}}$ are the numbers of datapoints strictly within the a and b marginal spaces of the k^{th}-nearest neighbour datapoints (including the i^{th} datapoint). For the example 2D density in Fig.6.3, $\mathrm{n_a}(i) = 6$ and $\mathrm{n}_b(i) = 4$.

k-nearest neighbours is able to estimate entropies with little bias because the numbers of datapoints within the k-nearest neighbour marginal spaces are allowed to vary for each datapoint $\mathrm{z}_i = (\mathrm{a}_i, \mathrm{b}_i)$. This allows the $\boldsymbol{\epsilon}$ to be the same in the joint and marginal spaces, meaning that the biases in the univariate and bivariate entropy estimates cancel. This can be seen by appreciating that Equ.6.17 is equivalent to Equ.6.16, but has not been simplified (including ϵ term cancellation).

[1] This expression for MI is equivalent to Equ.6.2 because of the conditional rule:

$$\mathrm{P(A|B)} = \frac{\mathrm{P(A,B)}}{\mathrm{P(B)}}. \tag{6.15}$$

$$\widehat{\mathrm{H}}(\mathrm{A}) + \widehat{\mathrm{H}}(\mathrm{B}) - \widehat{\mathrm{H}}(\mathrm{A,B}) = \frac{1}{\mathrm{N}} \sum_{i=1}^{\mathrm{N}} \log \epsilon(i) - \psi(\mathrm{n_a}(i)) + \psi(\mathrm{N}) + \frac{1}{\mathrm{N}} \sum_{i=1}^{\mathrm{N}} \log \epsilon(i) - \psi(\mathrm{n_b}(i)) + \psi(\mathrm{N}) - \frac{2}{\mathrm{N}} \sum_{i=1}^{\mathrm{N}} \log \epsilon(i) - \psi(\mathrm{k}) + \psi(\mathrm{N}). \quad (6.17)$$

6.3 Methods

To calculate the discriminabilities of early-stage neurodegenerative and healthy hybrid neuroimaging signals, k-nearest neighbours was used to estimate the MI between 1) disease-state label, **S**, and 2) sets of EEG DCM posterior parameters, **X**, and hemodynamics DCM posterior parameters, **Y**, at the four timepoints for which data were available: immediately before saline/LPS injection, 1 hour after, 2 hours after and 4.5 hours after. An intuitive interpretation of this MI is that it reflects the discriminability of the [**X**, **Y**] spaces between the two disease states. The more discriminable these spaces are, the more accurately they are likely to be classified (for clinical diagnosis) and the more significant their ANOVA statistics are likely to be (for disease monitoring and drug development).

To quantify the value of imaging neuronal activity and hemodynamics concurrently rather than separately, k-nearest neighbours was then used to estimate the synergy of the **S** and [**X**, **Y**] spaces. The intuitive interpretation of this synergy is that it is a measure of the relative discriminabilities of these spaces when the correlations between **X** and **Y**, i.e. aspects of neurovascular coupling, are captured compared to when they are not. A positive synergy results from instances where the discriminability of **S** is larger when the coupling of **X** and **Y** is included, supporting the development of hybrid neuroimaging systems for diagnosing, monitoring and understanding neurodegenerative diseases in their

earliest stages because the coupling of $\mathbf{X}$ and $\mathbf{Y}$ can only be inferred when the two imaging modalities are acquired concurrently.

The DCMs fitted to the neuroimaging data to generate $\mathbf{X}$ and $\mathbf{Y}$ are different from those fitted in the earlier chapters of this thesis because:

1. Single-modality DCMs are fitted to the electrophysiology and hemodynamics responses separately. This is because it is admissable that the dependencies between the parameters of Combined DCM are not accurate, and therefore if Combined DCM were fitted to the concurrent electrophysiology and hemodynamics responses, it could constrain the relationships of the multimodal parameters during fitting to not only be incorrect, but possibly biologically impossible[1]. By fitting separate DCMs to the electrophysiology and hemodynamics responses, no constraints are imposed on the relationships of the multimodal parameters in this chapter.

2. For the DCM fitting to the neuronal activity data, DCM-EEG could still be used (rather than DCM-LFP) because Rosengarten et al. recorded EEG responses (unlike the preliminary dataset, which contained LFP responses). The only difference that needed to be made to DCM-EEG was to allow parameter C (amplitude of thalamic neuronal activity arriving to the cortical column) to become negative by removing the exponentiation. This was necessary because there were clear differences in the polarities of the EEG responses not only between animals under the same condition, but also within animals (see Fig.6.4). These differences will have resulted from variations in the activation strengths of the different layers of the cortical column (the reader is reminded that the different cortical layers experience current sinks and sources in different orders, meaning that if one layer is

[1]It was less critical for the multimodal relationships to be accurate for the work carried out in Chapters 3 and 4 because the focus of these chapters was to investigate the consequences of the posterior parameter values on GLM significance, rather than to interrogate their relationships explicitly.

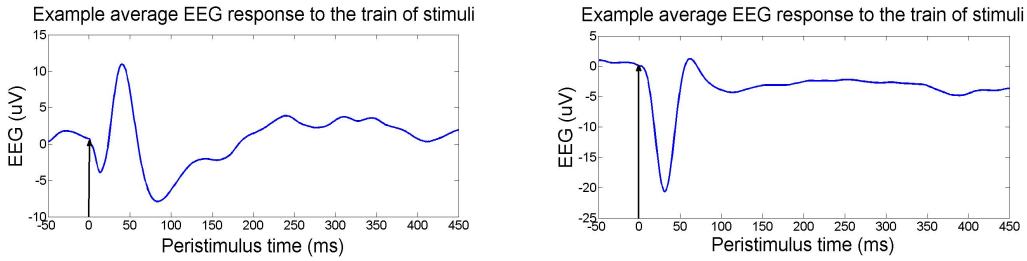

Figure 6.4: Two example EEG responses (averages over single trains of stimuli) recorded from different animals 2 hours after 5mg/kg LPS injection. Notice that in (a) the main EEG response is positive, suggesting that the EEG electrode recording this response was most sensitive to the more superficial cortical layers (layers I-III) (the majority of responses have this topology), whereas in (b) it is negative, suggesting that the EEG electrode recording this response was most sensitive to cortical layer IV. Other responses fall somewhere between these two extremes, suggesting that they reflect a mixture of these layers. Because of these differences in EEG responses across animals, it was necessary to free DCM-EEG parameter C to be negative so that all responses could be fitted to equally well.

activated more strongly than another, the EEG will reflect that layer's activity more, as explained in Fig.2.3).

3. For the DCM fitting to the hemodynamics data, it was necessary to modify mDCM-fMRI so that it models CBF only. This is because the hemodynamics recorded in Rosengarten et al.'s dataset are CBF rather than [ΔHbR]. This modified DCM is referred to as fDCM-fMRI and is shown in Fig.6.5.

6.3.1 Using k-nearest neighbours to estimate densities including categorical variables

Because disease-state label is not a random variable, it was necessary to derive a new formulation for k-nearest neighbours so that instead of, in the univariate case, H(X) + H(**S**) - H(X,**S**) having to be calculated, the equivalent H(X) - H(X|**S**) could be calculated (because it does not make sense to find k^{th}-nearest neighbour distances for a categorical variable such as disease-state label). To achieve this whilst ensuring that the biases in each entropy estimate still cancel, the k^{th}-nearest neighbour distances used in both the conditional, e.g. (X|**S**), and unconditional,

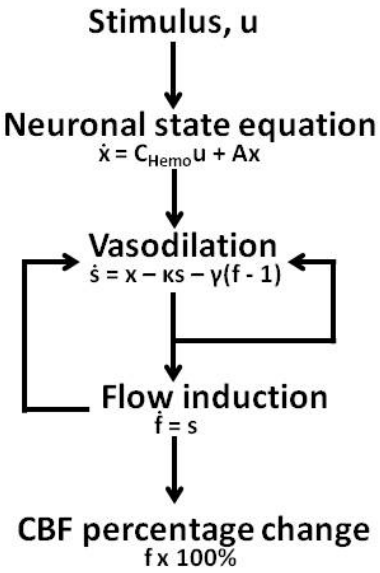

Figure 6.5: fDCM-fMRI for fitting to CBF percentage change data. Neural model free parameters, $\boldsymbol{\theta_N}$: A = self-inhibition of the electrical neuronal activity, x, and C_{Hemo} = amplification of the effect of the exogenous input, u, on x. Observer model free parameters, $\boldsymbol{\theta_O}$: κ = rate of vasodilatory signal decay; γ = rate of bloodflow autoregulatory feedback.

e.g. X, spaces need to be the same. In the original implementation of k-nearest neighbours, the length scales ($\boldsymbol{\epsilon}$) in the joint and marginal spaces were kept the same by making $\mathbf{n}_a$ and $\mathbf{n}_b$ free variables. A similar approach is adopted in order to include categorical variables (see Fig.6.6):

- A specified value for k is used when estimating the conditional entropies, as is done in the joint entropy estimation in the original k-nearest neighbours implentation. Call this $\mathrm{k}_{\mathrm{fixed}}$.
- Use the same k^{th}-nearest neighbour distances when calculating the unconditional entropies, as is done when estimating the marginal entropies in the original k-nearest neighbours implementation. This ensures that the distances used to estimate the conditional and unconditional spaces are the same, so that subtraction of one from the other results in their biases cancelling. However, using specified distances means that k is now the free variable, referred to as $\mathrm{k}_{\mathrm{free}}$ (analagous to n_a and n_b in the original

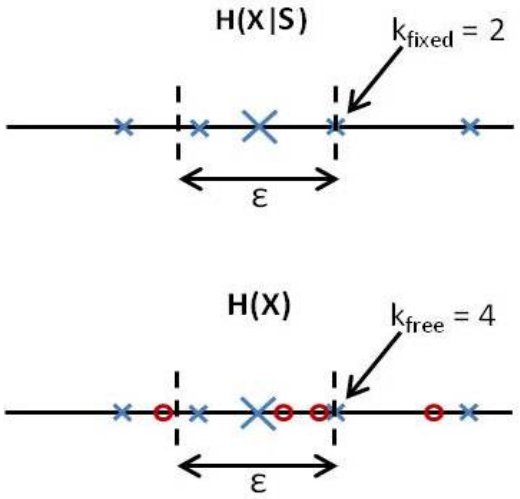

Figure 6.6: Diagrams showing the difference between k_{fixed} and k_{free} for the large blue cross datapoint (x_i) (1D example). When calculating H(X|S), i.e. the entropy of a DCM-EEG posterior parameter conditioned on diseased/healthy condition, the distance to the k^{th}_{fixed}-nearest neighbour datapoint, ϵ(i), is first found (in this example $k_{fixed} = 2$). This distance must then also be used when calculating H(X) (when the datapoints across both disease states are included), so that the biases in H(X) and H(X|S) cancel out. This can be achieved by setting k_{free} to be the number of datapoints across both disease states within $\frac{\epsilon(i)}{2}$ of datapoint i. In this example k_{free} would be 4.

implementation, but for the full rather than marginal spaces).

The MI between the DCM posterior parameter space [**X**, **Y**] and disease-state label **S** can now be estimated using:

$$\widehat{\mathrm{I}}([\mathbf{X},\mathbf{Y}];\,\mathbf{S}) = \frac{1}{\mathrm{N}}\sum_{i=1}^{\mathrm{N}} -\psi(\mathrm{k}_{free}(i)) + \psi(\mathrm{N}) + \psi(\mathrm{k}_{fixed}) - \sum_{s=1}^{\mathrm{S}} \mathrm{p(s)}\psi(\mathrm{N}_s), \qquad (6.18)$$

where N_s is the average number of datapoints in each disease state and p(s) is the probability of each disease state.

The equivalent algorithm for estimating the synergy of [**X**, **Y**] and **S**, i.e. the extra discriminability of disease-state label resulting from inclusion of different subsets of DCM posterior parameters [**X**,**Y**], is then:

$$\text{Synergy}([\mathbf{X},\mathbf{Y}];\,\mathbf{S}) = -\frac{1}{\mathrm{N}}\sum_{i=1}^{\mathrm{N}}\psi(\mathrm{k}_{free_{x,y}}(i)) + \frac{1}{\mathrm{N}}\sum_{i=1}^{\mathrm{N}}\psi(\mathrm{k}_{free_{x}}(i)) + \frac{1}{N}\sum_{i=1}^{\mathrm{N}}\psi(\mathrm{k}_{free_{y}}(i)) - \psi(\mathrm{N}) + \sum_{s=1}^{\mathrm{S}}\mathrm{p(s)}\psi(\mathrm{N}_s) - \psi(\mathrm{k}_{fixed}), \quad (6.19)$$

where $\mathbf{k}_{\mathbf{free}_x}$ and $\mathbf{k}_{\mathbf{free}_y}$ are the numbers of datapoints within the unconditional marginal spaces, and $\mathbf{k}_{\mathbf{free}_{x,y}}$ are the numbers of datapoints within the unconditional state joint space.

6.3.2 Finding optimal values for $\mathbf{k}_{\mathbf{fixed}}$

To identify the optimal value of $\mathrm{k}_{\mathrm{fixed}}$ to use for this dataset, the disease-state labels of the 4.5 hour DCM posterior parameter estimates were shuffled and k-nearest neighbours used to estimate their MI with disease-state label at a range of $\mathrm{k}_{\mathrm{fixed}}$ values (the disease-state label shuffling should theoretically cause the MI to be 0, meaning that any departure from 0 is caused by sampling error and bias). Kraskov et al. recommend that $\mathrm{k}_{\mathrm{fixed}}$ be tested between 2 to 8, and from their observations it would seem that $\mathrm{k}_{\mathrm{fixed}} = 8$ would be most appropriate for this dataset because they find that when the data length and correlation are low the optimal value for k approaches 8 [91]. This describes the neurodegenerative disease dataset because:

- The computational expense of DCM fitting limits the number of datapoints (DCM posterior parameter mean estimates) to a relatively low number (between 110 and 121 per disease state).
- Preliminary work indicated that the correlation of the DCM posterior parameter estimates on disease state was small.

However, because of the added complication of making k the free parameter when calculating H(X), which is calculated from approximately twice as much data as H(X|Y) (because there are two disease states), it is likely that k_{fixed} is optimal when less than 8. This is because if $k_{fixed} = 8$ then the average k_{free} will be 16, which will result in estimates for H(X) that are very insensitive to the data.

6.4 Results

The average EEG and CBF responses to a single train of stimuli for the saline and 5mg/kg LPS animals immediately before saline/LPS injection, 1 hour later, 2 hours later and 4.5 hours later are shown in Figs. 6.7 and 6.8, respectively. The most striking thing to notice from these figures is that there are clear differences between the saline and LPS EEG and CBF responses immediately before the saline/LPS injections, i.e. when no differences should exist between the groups. A two-sample t-test of the individual CBF percentage changes in the two groups found these differences to be extremely significant ($p < 0.001$). There is no obvious reason for these significant differences and so it is put down to subject variability.

These differences in the baseline EEG and CBF responses between the experimental groups are likely to be responsible for the responses not following the same pattern as Rosengarten et al.'s original dataset or the preliminary dataset from Chapter 5 at the later timepoints. For example, the mean CBF percentage changes 1 and 2 hours after injection are larger for the LPS animals than for the saline animals, rather than considerably smaller. Similarly, even though the N1-P1 EEG amplitude and CBF percentage change are reduced to the same extent 4.5 hours after saline/LPS injection in Rosegarten et al.'s original dataset, in the new dataset the CBF percentage change is only reduced to 65% of its baseline level whereas the N1-P1 EEG amplitude is reduced to 48% of its baseline level. This greater impairment of the neuronal activity than the hemodynamics could conceivably correspond to advanced neurodegenerative disease, where neuronal

death causes more neuronal death to occur, leading to greater reductions in neuronal activity than hemodynamics, or it could also be an artifact of the higher baseline level of CBF percentage change for the LPS group.

Unfortunately, these baseline differences complicate interpretation of the MI and synergy values. However, by calculating the MI and synergy of the data recorded immediately before injection, it is possible to comment on what inter-group differences occurring at the later timepoints may be caused by the baseline differences between the groups rather than by the LPS modulation.

Fig.6.10 and Fig.6.11 display the MI and synergy values calculated for all pairs of multimodal DCM posterior parameters using $k_{\text{fixed}} = 5$ (the optimal value for k_{fixed} for this dataset - see Fig.6.9(a)). The pairs, triplets and quadruplets of multimodal DCM parameters with the highest MI at each timepoint are listed in Tables 6.1 - 6.3, and the pairs, triplets and quadruplets with the highest and lowest synergies at each timepoint are listed in Tables 6.4 - 6.6.

6.5 Discussion

The significant differences in the baseline EEG and CBF responses between the experimental groups complicate interpretation of the MI and synergy values calculated in this chapter. However, one advantage of the new dataset over Rosengaten et al.'s original dataset is that the relative amplitudes of the average LPS and saline EEG responses decrease steadily with time since saline/LPS injection: the mean N1-P1 LPS amplitude is 101% that of the saline amplitude before baseline, 79% 1 hour after injection, 69% 2 hours after injection, and 48% 4.5 hours after injection. This suggests that the 1 hour data from this new dataset more closely corresponds to early-stage neurodegenerative disease than the data from Rosengarten et al.'s original dataset (where the reductions in neuronal activity were in fact slightly greater 1 hour after saline/LPS injection than 2 hours after). In this way, the new dataset provides insight for three stages of

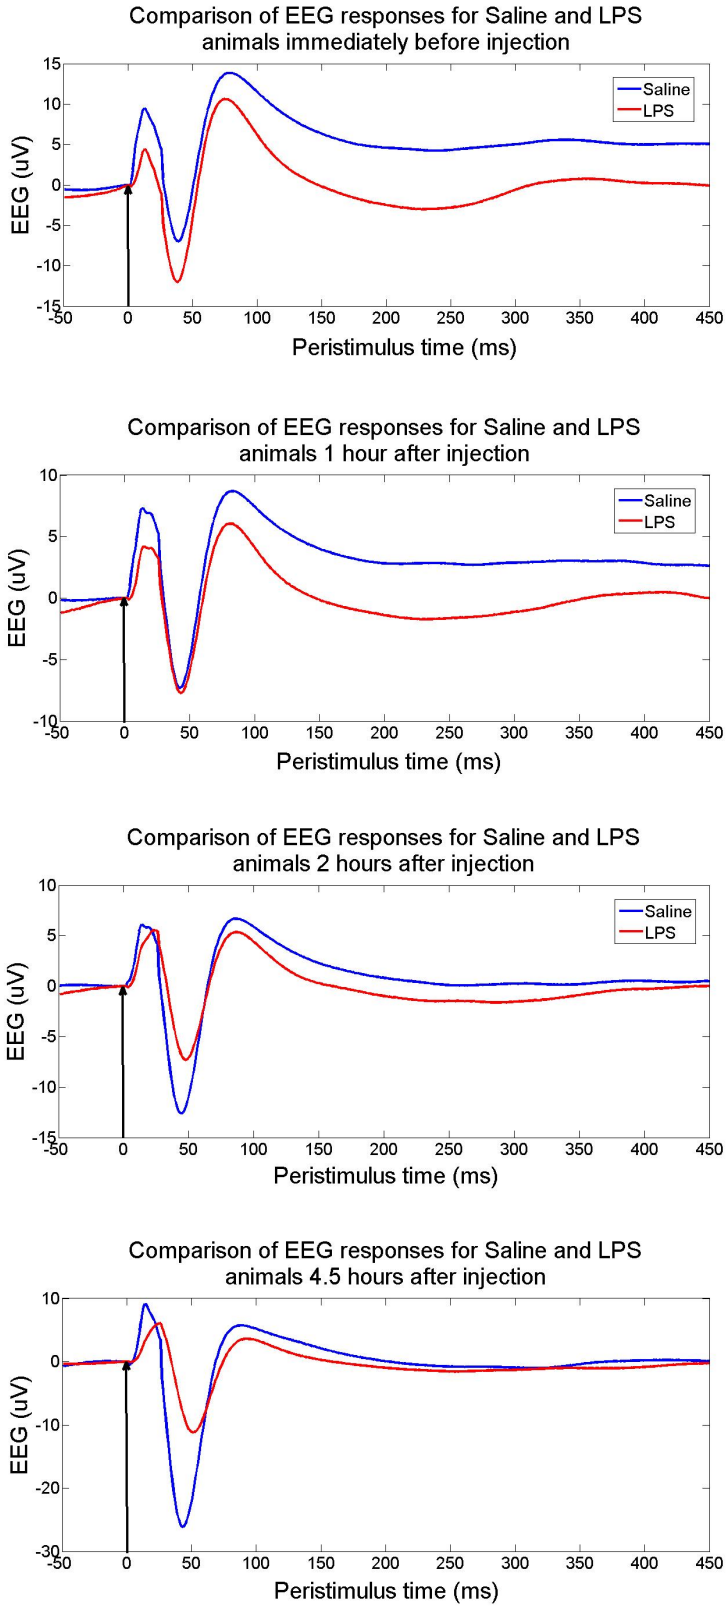

Figure 6.7: Comparison of average EEG responses to a single stimulus (a) immediately before, (b) 1 hour after, (c) 2 hours after and (d) 4.5 hours after saline or 5mg/kg LPS injection.

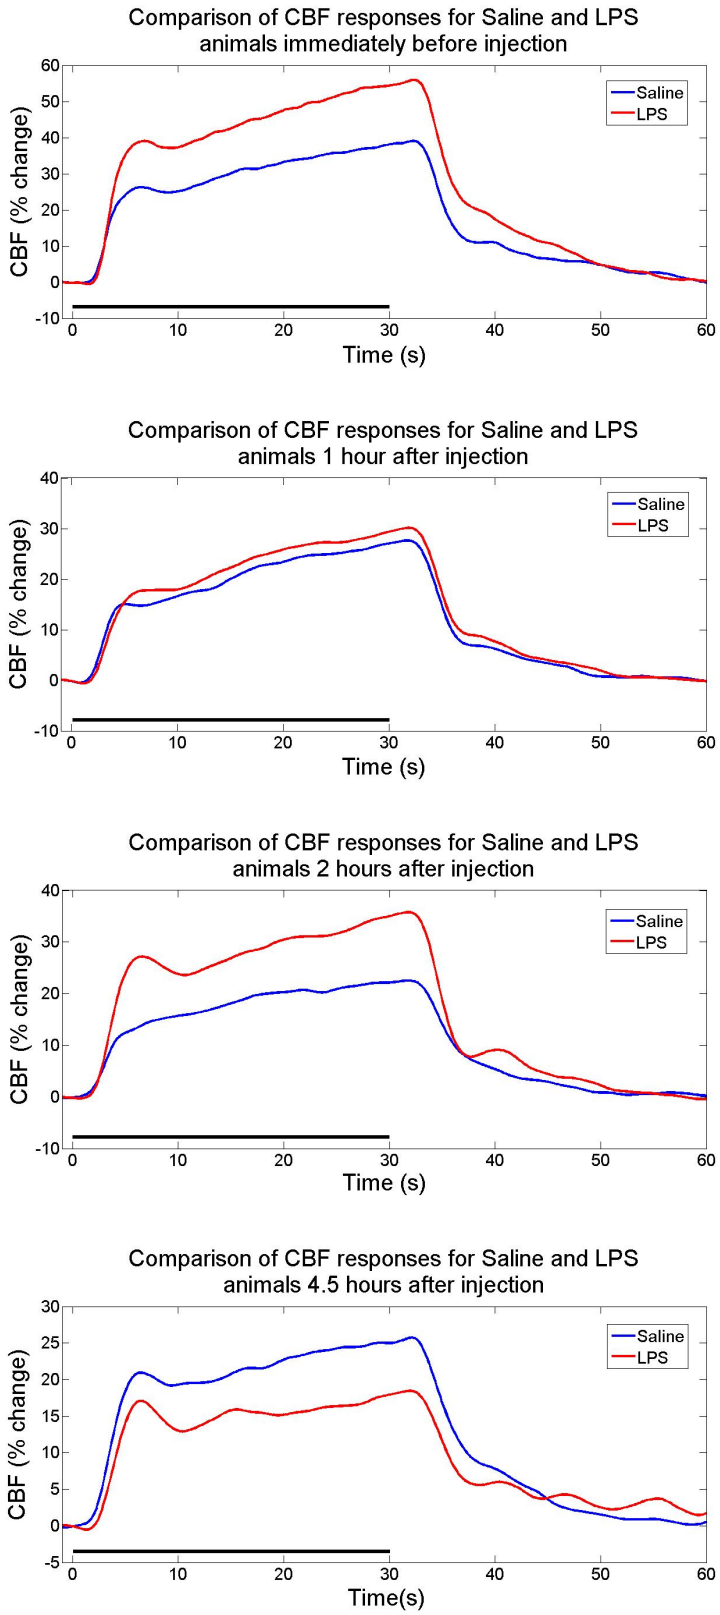

Figure 6.8: Comparison of average CBF percentage change responses to a 30s train of stimuli (a) immediately before, (b) 1 hour after, (c) 2 hours after and (d) 4.5 hours after saline or 5mg/kg LPS injection.

Rank	Beginning		1 hour		2 hours		4.5 hours	
	MI	**Parameters**	**MI**	**Parameters**	**MI**	**Parameters**	**MI**	**Parameters**
1	0.42	$\gamma 4$, γ	0.10	C, γ	0.10	$\gamma 4$, C_{Hemo}	0.28	$\gamma 4$, γ
2	0.41	$\gamma 4$, κ	0.09	C, κ	0.09	$\gamma 4$, A	0.28	$\gamma 4$, A
3	0.41	$\gamma 4$, A	0.09	C, A	0.08	H_i, A	0.27	$\gamma 4$, C_{Hemo}
4	0.40	$\gamma 4$, C_{Hemo}	0.09	C, C_{Hemo}	0.08	$\gamma 4$, γ	0.26	κ, γ
5	0.15	C, γ	0.08	$\gamma 4$, κ	0.08	H_i, κ	0.17	$\gamma 4$, A

Table 6.1: Pairs of DCM parameters with the five highest MI at each timepoint.

Rank	Beginning		1 hour		2 hours		4.5 hours	
	MI	**Parameters**	**MI**	**Parameters**	**MI**	**Parameters**	**MI**	**Parameters**
1	0.47	C, $\gamma 4$, γ	0.17	C, $\gamma 4$, κ	0.16	$\gamma 4$, A, C_{Hemo}	0.35	C, $\gamma 4$, γ
2	0.46	C, $\gamma 4$, A	0.17	C, H_i, κ	0.16	C, $\gamma 4$, γ	0.28	C, $\gamma 4$, C_{Hemo}
3	0.45	R1, $\gamma 4$, A	0.16	C, $\gamma 4$, C_{Hemo}	0.15	C, $\gamma 4$, C_{Hemo}	0.27	C, $\gamma 4$, A
4	0.44	R1, $\gamma 4$, γ	0.15	C, $\gamma 4$, γ	0.14	H_i, A, C_{Hemo}	0.26	R1, $\gamma 4$, C_{Hemo}
5	0.43	C, $\gamma 4$, A	0.15	C, H_i, κ	0.13	H_i, A, $\gamma 4$	0.17	$\gamma 4$, C_{Hemo}, γ

Table 6.2: Triplets of DCM parameters with the five highest MI at each timepoint.

Rank	Beginning		1 hour		2 hours		4.5 hours	
	MI	**Parameters**	**MI**	**Parameters**	**MI**	**Parameters**	**MI**	**Parameters**
1	0.47	R1, C, $\gamma 4$, γ	0.20	C, H_i, C_{Hemo}, γ	0.18	H_i, $\gamma 4$, A, C_{Hemo}	0.36	R1, $\gamma 4$, C_{Hemo}, γ
2	0.47	R1, C, $\gamma 4$, A	0.19	C, $\gamma 4$, C_{Hemo}, γ	0.17	$\gamma 1$, $\gamma 4$, A, C_{Hemo}	0.35	H_e, $\gamma 4$, C_{Hemo}, γ
3	0.47	C, H_i, $\gamma 4$, γ	0.19	C, H_i, $\gamma 4$, κ	0.17	τ_i, $\gamma 4$, A, C_{Hemo}	0.35	R1, $\gamma 4$, S2, C_{Hemo}
4	0.47	C, $\gamma 4$, S2, γ	0.18	C, H_i, C_{Hemo}, κ	0.16	τ_e, $\gamma 4$, A, C_{Hemo}	0.35	C, $\gamma 3$, $\gamma 4$, γ
5	0.47	C, H_e, $\gamma 4$, γ	0.18	C, $\gamma 4$, A, C_{Hemo}	0.16	C, H_e, $\gamma 4$, γ	0.35	C, $\gamma 4$, S2, γ

Table 6.3: Quadruplets of DCM parameters with the five highest MI at each timepoint.

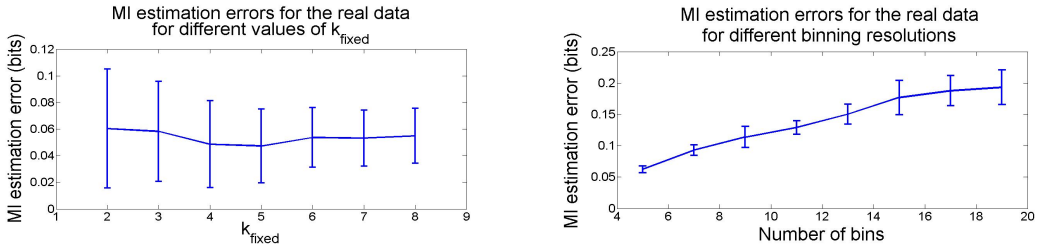

(a) (b)

Figure 6.9: MI estimation errors for (a) different k_{fixed} and (b) different binning resolutions. Notice that the relationship between k_{fixed} and MI estimation error is very weak, especially compared to the relationship between number of bins and MI estimation error, supporting the claim that k-nearest neighbours is less sensitive to unoptimal choice of its free parameter (k_{fixed}) than other MI estimators [147]. As the mean error is smallest when $k_{fixed} = 5$, this value was chosen for all information-theoretic analyses in this chapter.

neurodegenerative disease progression: the 1 hour dataset most closely corresponds to early-stage neurodegenerative disease, the 2 hour dataset most closely corresponds to established neurodegenerative disease, and the 4.5 hour dataset most closely corresponds to advanced neurodegenerative disease. Although there is less value in being able to diagnose and understand the later stages of neurodegenerative disease, as little effective treatment is available to these patients, it is nevertheless interesting to understand whether generative embedding is able to discriminate early-stage and advanced neurodegenerative disease states. This is because it goes someway towards understanding the potential of generative embedding for sub-typing neurodegenerative diseases[1].

6.5.1 MI

The first thing to notice from Fig.6.10 is that the baseline MI for each of the multimodal pairs including $\gamma 4$ is higher than for any other pairs of multimodal parameters at any timepoint. This suggests that $\gamma 4$ is not only strongly responsible for the clear differences in the average EEG responses between the two

[1]Brodersen et al. recently published a study reporting that they were able to classify DCM-fMRI visual-parietal-prefrontal connection strengths into one of three subgroups of schizophrenia with 78% accuracy [17]

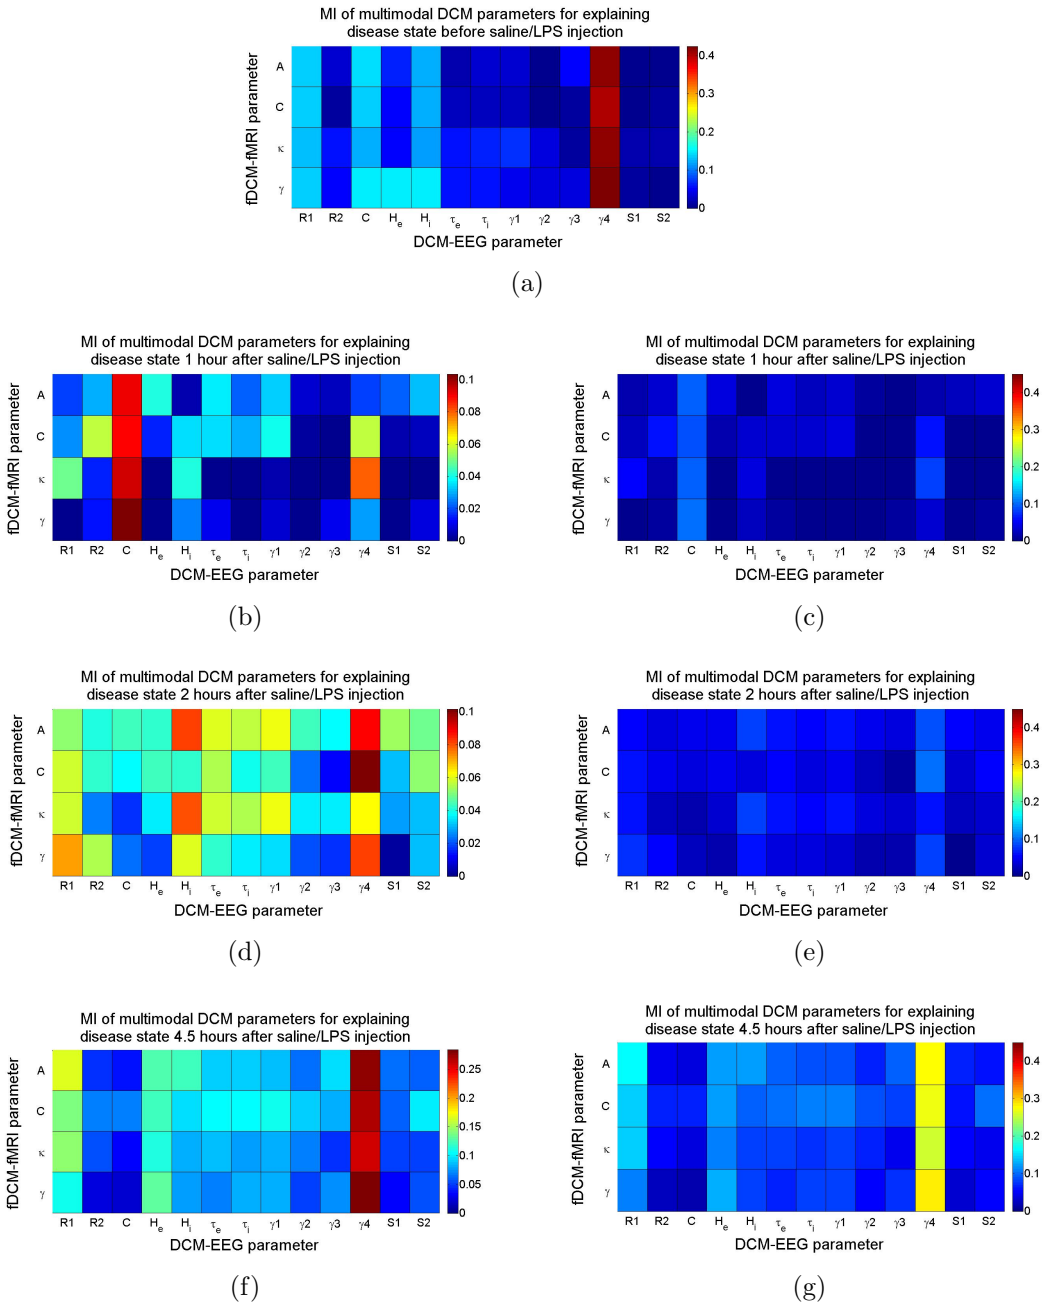

Figure 6.10: MI of pairs of multimodal DCM parameters for explaining disease state (saline or LPS group label) estimated using k-nearest neighbours (a) immediately before, (b) and (c) 1 hour after, (d) and (e) 2 hours after, and (f) and (g) 4.5 hours after saline or 5mg/kg LPS injection. (a), (b), (d) and (f) have separate colour axes, whereas the colour axes of (c), (e) and (g) match the colour axis for (a), which has the largest MI range, to enable comparison of MI across the timepoints.

groups before saline/LPS injection (see Fig.6.7), but also for the subsequent clear differences in the average CBF responses (see Fig.6.8).

The baseline MI of R1, C and H_i paired with any of the fDCM-fMRI parameters are also above average estimation error levels (approximately 0.05 bits for k_{fixed}, as shown in Fig.6.9(a))[1], suggesting that these neuronal activity processes also contribute to the baseline differences in the average EEG and subsequent CBF responses[2].

The second thing to notice from Fig.6.10 is that MI between $\gamma 4$ and each of the fDCM-fMRI parameters is considerably lower 1 and 2 hours after saline/LPS injection than at baseline. This suggests that the LPS either reduces or masks the natural variation in this parameter between the experimental groups. However, the MI between $\gamma 4$ and each of the fDCM-fMRI parameters is partly restored to baseline levels by 4.5 hours after saline/LPS injection, suggesting that by this timepoint either the naturally occurring $\gamma 4$ can once again be detected/expressed, or that increases in the discriminabilities of this parameter have occurred independent of the baseline differences, i.e. due to the LPS. If the latter case is true, then $\gamma 4$ paired with some combination of the fDCM-fMRI parameters should be the optimal parameter set for classifying/regressing the disease state simulated 4.5 hours after LPS injection, i.e. advanced neurodegenerative disease state.

The pattern of MI across the timepoints for R1 paired with any of the fDCM-fMRI parameters is similar to that just described for $\gamma 4$ (reduced at 1 and 2 hours after saline/LPS injection but then partially restored by 4.5 hours after saline/LPS injection). However, the pattern is different for the other two

[1]The average MI estimation error of 0.05 bits when $k_{fixed} = 5$ is smaller than the average error estimated for the same data but using binning at any binning resolution tested - see Fig.6.9(b).

[2]The high baseline MI have only been attributed to the neuronal activity processes because their MI with all four fDCM-fMRI parameters are consistently high. Had any of the four fDCM-fMRI parameters had consistently high MI with a number of DCM-EEG parameters, it would suggest that that hemodynamics process was responsible for the differences in the CBF responses between the groups instead. Alternatively, if specific multimodal pairs had had high MI without those parameters exhibiting high MI with other parameters, it would indicate that the coupling between those multimodal parameters was discriminable between the groups, rather than the DCM-EEG or fDCM-fMRI parameter on its own.

DCM-EEG parameters that have high baseline MI with all of the fDCM-fMRI parameters: C and H_i have relatively high MI 1 and 2 hours after the saline/LPS injection, respectively, but relatively low MI 4.5 hours after saline/LPS injection. These are the only two sets of parameter pairs to follow this pattern, as generally the MI of all of the multimodal parameter pairs increases from 1 hour to 2 hours to 4.5 hours after saline/LPS injection. Therefore, ignoring the parameters with high baseline MI (γ4, R1, C and H_i), and assuming that time since LPS injection relates to neurodegenerative disease progression, the discriminabilities of the multimodal parameter spaces generally increases with neurodegenative disease progression. This finding agrees with the first prediction made in Chapter 5 that disease-state discriminability would increase with disease progression if negative multimodal couplings could be captured.

Unfortunately, because some parameters have high baseline MI, it is not useful to discuss which specific sets of multimodal DCM parameters ought to be used in classification or ANOVA statistics, or even confirm whether the spaces with the highest MI have negative multimodal couplings, as predicted. However, it is useful to see that, apart from the data acquired before saline/LPS injection, which has MI that plateaus at 0.47 bits for the multimodal triplets[1], the MI for the multimodal triplets are higher than for the pairs, and the MI for the multimodal quadruplets are very slightly higher than for the triplets. This indicates that there is value in classifying/regressing spaces of at least four multimodal parameters, but it is unlikely that the small increases in discriminabilities resulting from more than four parameters being included will justify the added computation time.

[1]The fact that the MI of the pre-injection parameters plateaus when only three parameters are included suggests that the inter-group differences between the baseline EEG and CBF responses are limited to a few, specific biological differences.

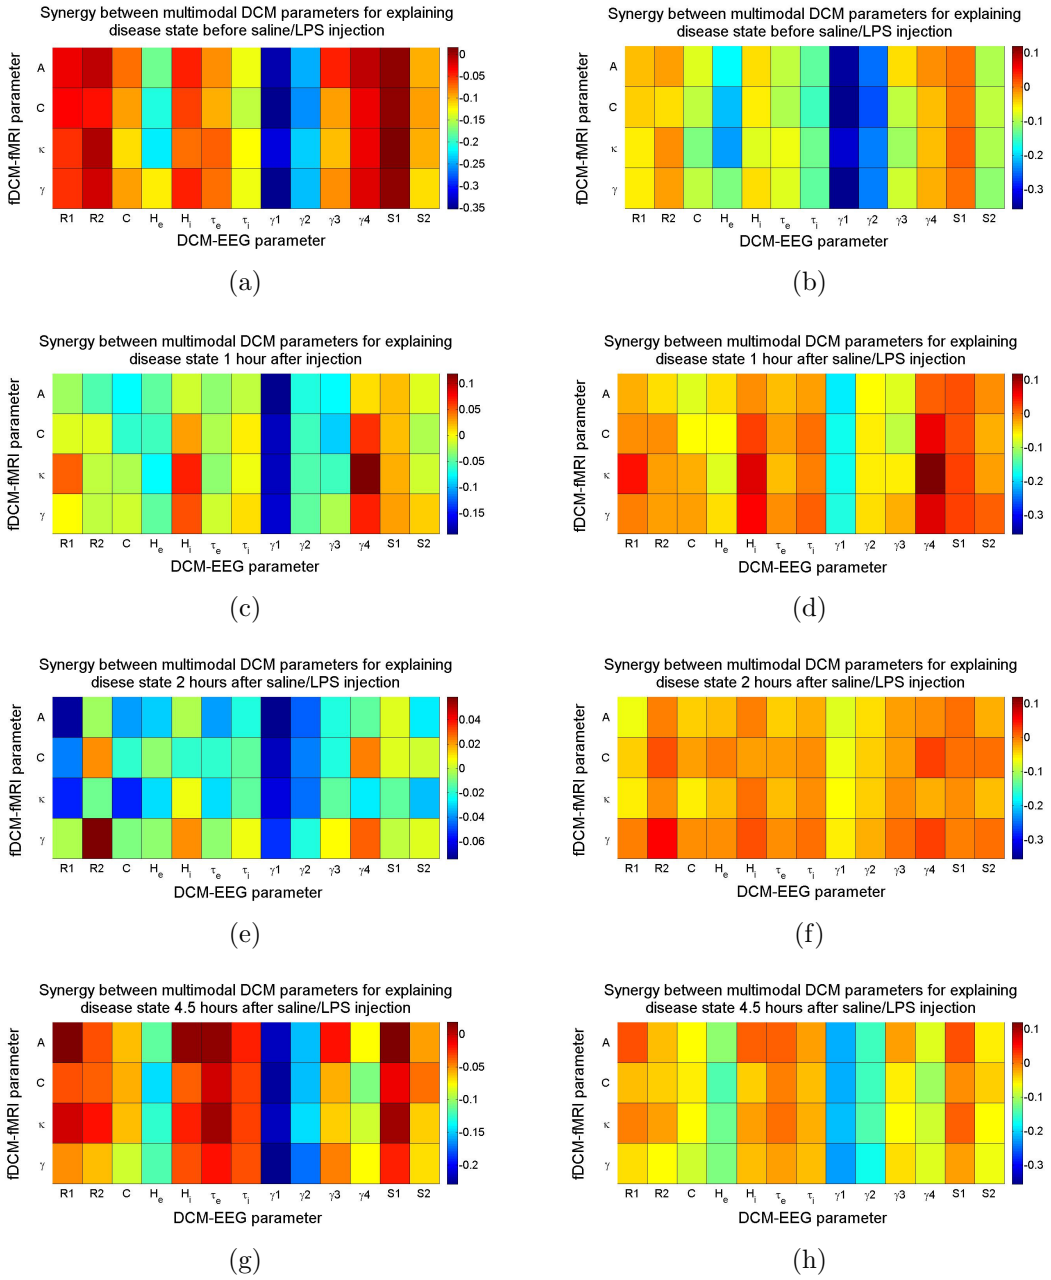

Figure 6.11: Synergy between pairs of multimodal DCM parameters for explaining disease-state (saline or LPS group label) estimated using k-nearest neighbours (a) and (b) immediately before, (c) and (d) 1 hour after, (e) and (f) 2 hours after, and (g) and (h) 4.5 hours after saline or 5mg/kg LPS injection. (a), (c), (e) and (g) have separate colour axes, whereas (b), (d), (f) and (h) have identical colour axes to enable comparison of synergies across the timepoints.

Rank	Beginning		1 hour		2 hours		4.5 hours	
	Synergy	**Parameters**	**Synergy**	**Parameters**	**Synergy**	**Parameters**	**Synergy**	**Parameters**
1	0.015	S1, κ	0.120	γ4, κ	0.059	R2, γ	0.018	γ4, A
2	0.009	S1, C_{Hemo}	0.070	γ4, γ	0.029	γ4, γ	0.017	R1, A
3	0.009	S1, A	0.068	H_i, κ	0.08	H_i, A	0.014	H_i, A
4	0.007	S1, γ	0.064	γ4, C_{Hemo}	0.022	H_i, γ	0.011	τ_e, A
5	-0.007	R2, κ	0.053	H_i, γ	0.022	R2, C_{Hemo}	0.009	S1, κ

Table 6.4: Pairs of DCM parameters with the five highest synergies at each timepoint.

Rank	Beginning		1 hour		2 hours		4.5 hours	
	Synergy	**Parameters**	**Synergy**	**Parameters**	**Synergy**	**Parameters**	**Synergy**	**Parameters**
1	0.019	S1, A, γ	0.125	γ1, γ4, κ	0.077	R2, γ, γ3	0.017	R1, γ1, A
2	0.018	S1, C_{Hemo}, γ	0.121	R1, γ4, κ	0.065	R2, H_e, γ	0.017	H_i, S1, A
3	0.018	S1, C_{Hemo}, κ	0.120	τ_i, γ4, κ	0.065	R2, S1, γ	0.016	H_i, τ_i, A
4	0.018	S1, κ, γ	0.119	γ3, γ4, κ	0.063	R2, S2, γ	0.015	γ1, A, S1
5	0.017	S1, A, C_{Hemo}	0.116	γ2, γ4, κ	0.059	R2, A, γ	0.15	H_i, γ1, A

Table 6.5: Triplets of DCM parameters with the five highest synergies at each timepoint.

Rank	Beginning		1 hour		2 hours		4.5 hours	
	Synergy	**Parameters**	**Synergy**	**Parameters**	**Synergy**	**Parameters**	**Synergy**	**Parameters**
1	0.017	S1, A, κ, γ	0.125	γ1, γ2, γ4, κ	0.078	R2, τ_e, γ3, γ	0.030	C_{EEG}, τ_e, C_{Hemo}, κ
2	0.017	S1, A, C_{Hemo}, κ	0.125	τ_i, γ1, γ4, κ	0.078	R2, γ2, γ3, γ	0.027	C_{EEG}, τ_i, C_{Hemo}, κ
3	0.016	S1, C_{Hemo}, κ, γ	0.124	γ1, γ3, γ4, κ	0.075	R2, γ1, γ3, γ	0.026	C_{EEG}, γ1, C_{Hemo}, κ
4	0.015	S1, A, C_{Hemo}, γ	0.123	R1, γ1, γ4, κ	0.074	R2, γ3, A, γ	0.022	C_{EEG}, γ2, C_{Hemo}, κ
5	0.014	C_{EEG}, γ1, S1, A	0.122	γ3, γ4, S1, κ	0.074	R2, H_e, γ3, γ	0.018	H_i, τ_i, S1, A

Table 6.6: Quadruplets of DCM parameters with the five highest synergies at each timepoint.

6.5.2 Synergy

The interesting results from Fig.6.11 and Tables 6.4 - 6.6 concern the positive synergies only, as only the positive synergies lend support to hybrid imaging systems, i.e. concurrent imaging of neuronal activity and hemodynamics.

The main observation to make from these results is that the synergies between the multimodal parameter pairs, triplets and quadruplets generally increase from baseline to 1 hour, and then from 1 hour to 2 hours. However, those parameter sets that were particularly redundant at baseline (those including $\gamma 1$, $\gamma 2$ and H_e) decrease back towards baseline levels by 4.5 hours. This indicates that the value of concurrent imaging of neuronal activity and hemodynamics across all parameter sets increases with disease progression up to a point (the disease state corresponding to LPS modulation some time between 2 and 4.5 hours after injection), and then it decreases back towards baseline levels. Assuming once again that time since LPS injection relates to neurodegenerative disease progression, the value of hybrid imaging across all multimodal parameter sets appears to be larger for early-stage neurodegenerative disease (when hemodynamics are modulated but neuronal activity is relatively unmodulated) and established neurodegenerative disease (when both hemodynamics and neuronal activity are modulated) than for late-stage neurodegenerative disease (when hemodynamics are modulated and neuronal activity is very modulated). Given that the individual highest synergy is reported for parameter sets including $\gamma 4$ and κ 1 hour after saline/LPS injection, the main conclusion to be made from the synergy calculations is that early-stage disease diagnosis, monitoring and understanding would benefit most from the development of hybrid imaging systems supported by generative embedding analyses. This finding agrees with the second prediction made in Chapter 5, that the synergistic effect of neurovascular uncoupling will result in the discriminability of the disease states being largest for the early stages of neurodegenerative disease.

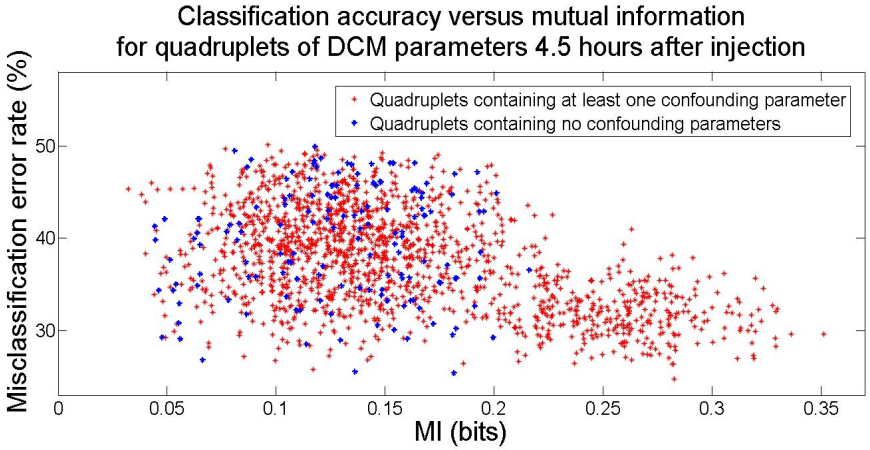

Figure 6.12: Plot showing how the classification accuracy of the quadruplets of DCM parameters recorded 4.5 hours after saline/LPS injection varies with MI.

6.5.3 Technological viability

The first main finding of this chapter was that the discriminabilities of subsets of multimodal DCM posterior parameters increase with disease progression. To gain some insight into how these discriminabilities translate into statistics for diagnosing neurodegerative disease (one of many possible application-specific generative embedding algorithms that this chapter informs), Matlab's "classify" function was used to fit multivariate normal densities with covariance estimates stratified by disease-state group (the "quadratic" algorithm type) to the quadruplets of DCM parameters recorded 4.5 hours after saline/LPS injection (the parameters found to have the highest MI values in the previous section). As can be seen in Fig.6.12, there does not appear to be a relationship between MI and classification accuracy for MI less than approximately 0.2 bits (although there does appear to be a relationship above this threshold), and unfortunately all of the quadruplets that do not contain a confounding parameter (the DCM parameters with high baseline MI: $\gamma 4$, R1, C, H_i) fall below this threshold. Nevertheless, the average misclassification rate below this threshold is 39% (dropping to less than 32% above it), and this is extremely significantly different from random classification.

These misclassification accuracies are higher than those reported by Fazli et al.

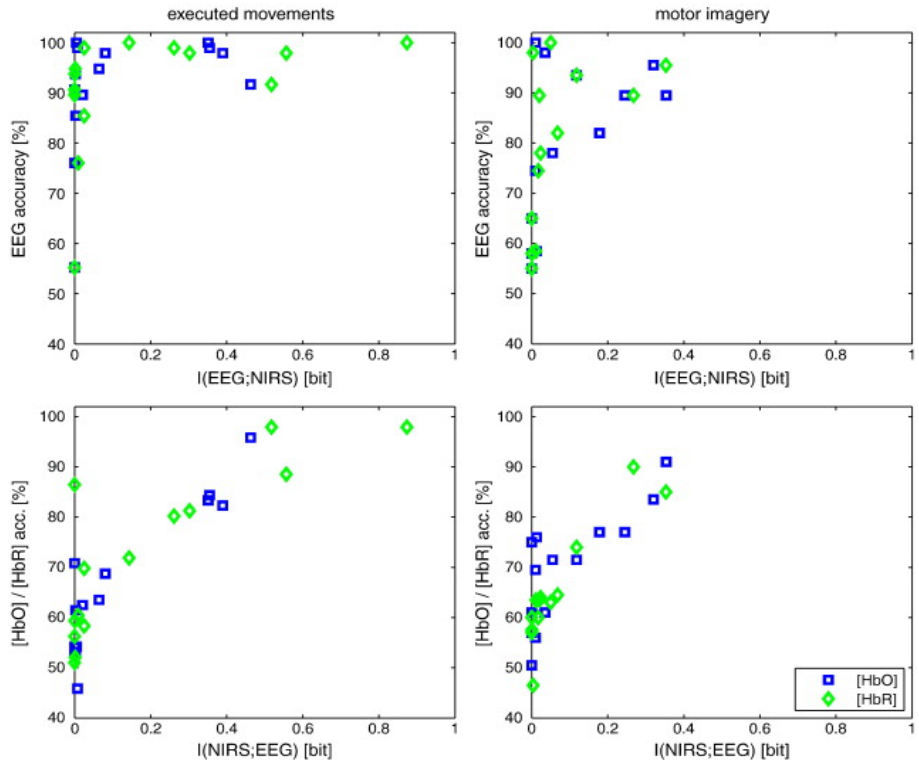

Figure 6.13: Classification accuracies of EEG and NIRS timeseries features plotted against their MI for explaining left or right hand motor execution/imagery. Each datapoint represents the results of a single subject (blue = [HbO]; green = [HbR]). Figure reproduced from Fazli et al., Figure 8 [52].

in their work comparing 1) the MI of hybrid EEG and near-infrared spectroscopy (NIRS, which measures blood oxygenation) feature spaces (datapoints from the EEG and NIRS response timeseries) for explaining left or right hand motor execution/imagery, against 2) classification accuracy of the EEG and NIRS features spaces separately - see Fig.6.13. There are numerous reason why this might be, such as classification of signals from motor cortex being easier than classification of signals from somatosensory cortex, but it nevertheless suggests that the potential of generative embedding for diagnosing neurodegenerative disease state (and perhaps even subtype) is greater than it would seem from the results presented in this chapter.

In summary, the findings made in this chapter lend support to the hypothesis that using concurrent imaging of neuronal activity and hemodynamics to diagnose, monitor and understand early-stage neurodegenerative diseases is not only a viable

practice for both clinical use and drug development, but that the greatest value of developing hybrid systems is for the early stages of neurodegenerative disease, rather than for the later stages.

Chapter 7

Commercialising generative embedding for investigating neurovascular coupling

Research and technology in this field, for either clinical, basic research or commercial application, is costly. A major driver in developing and applying the early-stage technical research presented in this thesis is therefore commercialisation, as a return on investment can more easily be seen. This chapter explores the commercial potential of the generative embedding work presented in this thesis as a possible route for maximising its impact, culminating in a set of recommendations for achieving this commercialisation.

7.1 Introduction

In 2004, the US Food and Drugs Administration (FDA) published a report in which they exposed a slowdown in innovative medical therapies reaching patients, despite exponential increase in basic science advancements. They attributed this slowdown to the increasingly challenging, inefficient and costly section of the product development pathway that is concerned with establishing the safety,

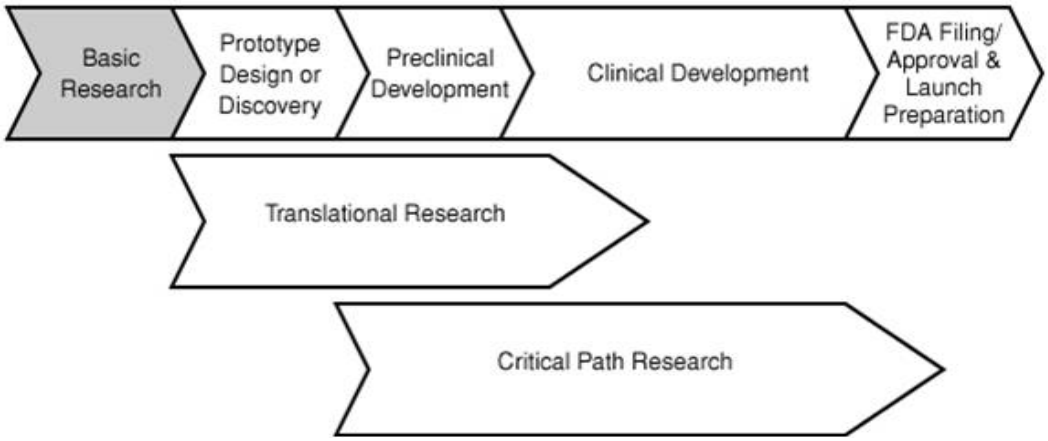

Figure 7.1: Medical therapy product development path. Reproduction of Figure 5 from [132].

medical utility and industrialisation of translational research, namely critical path research (see Fig.7.1) [132]. There is a great need to address this decrease because the direct impact that medical devices have on healthcare is clear: in a poll of primary care physicians, magnetic resonance imaging (MRI) and computer-tomography (CT) scanning were identified as the medical innovations that had had the greatest impact on the care they had provided over the previous 25 years [63]. What is more, medical devices can also have a substantial impact on healthcare indirectly: as the innovation rate of pharmaceutical companies is not matching the 20-year patent life of their blockbuster drugs (see Fig.7.2), pharmaceuticals are being forced to revolutionise their critical path research. The potential for new medical devices to play a role in this revolution of drug development is huge.

7.2 Potential commercial products

During this DPhil, generative embedding algorithms have been developed that provide insight into neurovascular coupling. These algorithms have the potential to form the basis of commercialisable software that aids in the diagnosis, monitoring and treatment of neurodegenerative patients imaged using innovative hybrid

Figure 7.2: The recent patent cliff is evidence that pharmaceutical companies are currently not innovating fast enough to be sustainable. Reproduced from [194].

imaging solutions. There are two obvious markets that could benefit from this potential: clinical healthcare and drug development.

The potential novel value that this software could bring to the clinic arises from the fact that neurovascular coupling could act as a biomarker that is routinely used to diagnose and monitor patients with neurodegenerative disease at earlier timepoints in disease progression and with greater accuracy than is presently achievable.

The first step of the generative embedding algorithms in this software product would be the fitting of biophysical models similar to DCM-EEG and DCM-fMRI to concurrently acquired neuronal and hemodynamic signals. In order to diagnose diseases caused/accompanied by impaired neurovascular coupling (such as early-stage neurodegenative disease), the resulting multimodal posterior parameters would be embedded into classification algorithms that attempt to classify the multimodal posterior parameter space of the new patient using a reference database of patients with known disease status. If the aim were instead to monitor the progression of neurovascular coupling (either deterioration or improvement as a result of treatment), the multimodal posterior parameter

estimates could be embedded into ANOVA statistics to detect whether specific neurobiological processes have significantly changed between timepoints.

Because neurodegenerative diseases are not currently diagnosed or monitored using imaging (or in fact any quantitative method), using imaging to diagnose neurodegenerative diseases represents radical innovation for the clinic.

Pharmaceutical companies could also benefit from this novel method for detecting whether specific neurobiological processes have significantly changed between timepoints because it would provide them with a new source of evidence for the efficacy and safety of their candidate neurodegenerative disease drugs. As some imaging techniques are already part of the drug development toolkit, using imaging to investigate aspects of neurovascular coupling represents incremental innovation for the industry.

7.2.1 Imaging modalities

To realise which modalities for imaging neuronal activity and neural hemodynamics would need to be combined to create the optimal hybrid imaging solutions for the clinical healthcare and pharmaceutical markets, it is necessary to compare their individual properties. Table 7.1 compares EEG and magnetoencephalography (MEG), which are the two modalities available for non-invasive imaging of neuronal activity, as well as near infra-red spectroscopy (NIRS) and fMRI, which are the two modalities available for non-invasive imaging of hemodynamics.

It is possible to integrate all of these modalities, and so the optimal hybrid system for concurrent imaging of neuronal activity and hemodynamics depends on the market in question.

	Spatial localisation	Temporal resolution	Sensitivity	Cellular- or systems-level	Capital and running costs	Portability	Motion artifacts	Simplicity
EEG	very low	very high	low	systems	low	very high	low	high
MEG	low	very high	very low	systems	very high	none	medium	very low
NIRS	very low	very high	medium	cellular	very low	very high	low	very high
fMRI	very high	very low	medium	systems	high	none	very high	very low

Table 7.1: Properties of non-invasive imaging modalities of neuronal activity (EEG and MEG), and neural hemodynamics (NIRS and fMRI). The quality of the information available in the signals is a function of their spatial localisation, temporal resolution and sensitivity. Modalities sensitive to motion artifacts require the patient to be very still during imaging, which is not possible for some patient populations. Simplicity refers to simplicity to operate and repair the hardware (and any associated software).

7.3 Markets

The global aging epidemic is unprecedented, pervasive (has global implications) and enduring, and therefore the market for ageing-related healthcare will continue to increase for a long time [130]. In addition to this, the number of medical imaging procedures performed is growing. This growth is fuelled by [181]:

- baby boomers demanding more services,
- technological advancements (including introduction of less invasive procedures),
- consumerism (greater demand and expectations of greater speed and quality).

For these reasons, there is a clearly growing market for products that directly improve ageing-related healthcare, including imaging tools that aid in the diagnosis, monitoring and treatment of neurodegenerative patients.

There is also a clear need to improve the way drugs for neurological disorders are developed. Because the brain is by far the most complex and least understood organ in the body, the success rate of neurological drugs making it through clinical trials to public health agency approval is much lower than the (already very low) average across all drugs, and as a consequence a number of pharmaceutical companies are currently withdrawing their investment from research into these conditions [20].

It is expected that neuroscience research improving our ability to identify and predict the action of novel neurological drugs will be greatly accelerated over the next decade. This is because the overarching aims of the publicly funded US Brain Initiative and EU Human Brain Project, which will both run for ten years, are to gain a comprehensive understanding of the whole brain at the cellular level. This will make the neurological drug market much more attractive for pharmaceutical companies. However, in order to be competitive in this future market, it will be necessary for them to invest in innovations, such as neuroimaging tools, that make

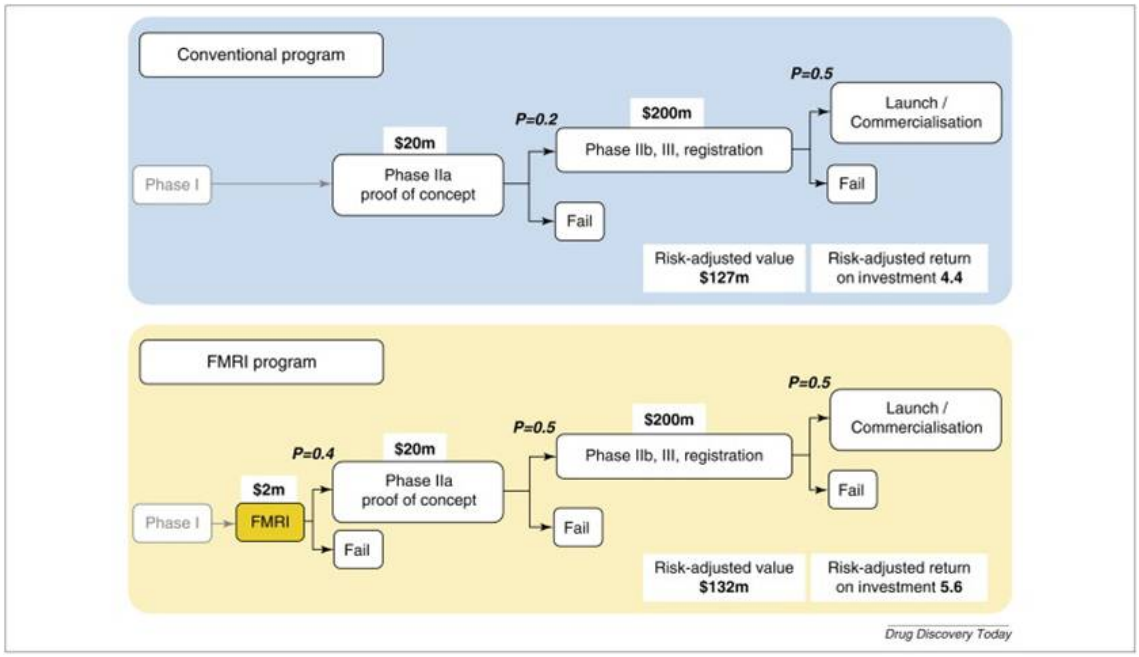

Figure 7.3: Costs predicted by a pharmaceutical company for getting a drug through clinical trials (top) with and (bottom) without using fMRI. Figure reproduced from Wise et al, 2010, Figure 3 [196].

their neurological drug clinical trials pipeline as efficient as possible. Fig.7.3 explains the theoretical financial advantage that fMRI could have for neurological drug development.

When considering which imaging modalities are most likely to gain traction in these clinical and drug development markets, it is important to consider their current usage as well as the specific requirements of these markets.

7.3.1 Imaging in the clinic today

Both EEG and MEG are currently used in the clinic for investigating and monitoring seizure disorders (such as epilepsy), sleep disorders, head injuries, brain tumours, coma and brain death. Although some district general hospitals in the UK still do not have an on-site EEG machine or appropriately trained staff [179], EEG is considered an established clinical tool. On the other hand, MEG has only very recently entered the clinical market, and it is currently only available in

Country	2006-7	2008-9	2010-11
Botswana	0	0	0.5
Ghana	0	0	0.16
Greece	16.3	19.6	22.6
Iceland	19.7	18.8	22.0
Japan	40.1	43.1	-
Mexico	1.4	1.7	2.0
Poland	1.9	2.9	4.7
UK	3.6	7.3	9.5
USA	5.6	5.5	5.9

Table 7.2: Number of MR scanners per million population in a selection of countries [54][152][35].

highly specialised hospitals [185].

NIRS is the technology behind a number of devices that currently have widespread use in the clinic, including the more than 2,000 cerebral oximeters available today for making crude measurements of cerebral oxygenation during queries of neural ischemia, as well as the instrument used to image the retina when assessing macular degeneration. Neural NIRS imagers with the capability to image hemodynamics in detail are only beginning to transition from research to the clinical market, with only a handful of devices having received FDA approval at this time [197].

MR scanners are available in hospitals around the world to diagnose and understand the extent of diseases or trauma throughout the body. They have widespread use in developed countries (see Table 7.2), however only 5-10% of the world's MR scanners are sufficiently powerful enough to measure brain hemodynamics using fMRI, i.e. that have a magnet strength of 3T or higher [48]. With the recommended lifetime of MR scanners being seven years [129], and because the capital cost of increasingly powerful magnets is more or less constant, it is likely that these faster and more powerful magnets will be commonplace in hospitals within the next ten years.

fMRI currently only sees routine use in a select number of hospitals and as a

tool for presurgical planning, i.e. to assist with mapping of cortical areas in order to minimise impact on non-diseased functional areas during neurosurgery [20]. The pioneering hospitals that are already using fMRI for presurgical planning, such as Great Ormond Street in the UK [78], use the software that comes as standard with MR scanners, i.e. developed by Siemens (syngo.MR Neuro), General Electric (BrainWave) or Philips (IView BOLD).

Hybrid EEG-fMRI[1] is a very nascent clinical technology for locating epileptogenic zones that is used in a small number of pioneering hospitals, such as Inselspital Bern and Hopitaux Universities de Geneve in Switzerland [165].

7.3.2 Imaging in drug development today

The aim of a clinical trial is to gather evidence that a novel compound is both safe and effective so that regulatory bodies such as the FDA and the National Institute for Health and Clinical Excellence (NICE, the UK's public health agency) approve it as a medical therapy. Positron emission tomography (PET) is currently the standard tool used to understand the cellular pathologies of diseases, to visualise drug-target engagement, and to measure pharmacokinetics and pharamacodynamics. However, PET imaging suffers a number of drawbacks:

- PET can only image cellular-level processes, but systems-level markers of drug action are:
  - more indicative of clinical therapeutic effect, meaning they are more likely to provide clinical endpoints of drug success,
  - are more translatable across species, which is important because preclinical work occurs in the beginning stages of clinical trials, and so if the safety and efficacy of novel drugs can be better predicted at this

[1] EEG-fMRI is made possible by using unmagnetic materials to measure the EEG (to ensure patient safety and reduce artifacts in the fMRI signal) and post-processing to remove artifacts from the EEG signals resulting from the fast changes of the MR magnetic field.

early stage, less money will be spent bringing drugs that ultimately fail through the more costly human trials [15],

  - useful for stratifying subjects, improving the rates that drugs meet their clinical endpoints because the trial will be run on fewer non-responders.

- PET is the imaging of radioactive substances in the body, and limits on the amount of radiation that it is safe to expose a subject to make it difficult to collect longitudinal data.
- PET ligands need to be developed for every new target, and this is not always achieved.
- PET is costly, especially because in order to use ligands with short half-lives (so that the radiation decays away faster), the pharmaceutical company needs to create their own isotopes, which means they must own and operate a cyclotron.

Over the last decade, EEG has also seen some routine use in drug development, however, because standards have not been defined for its use as a biomarker of neurological drug efficacy, it is only currently routinely used to establish the safety of novel compounds on the central nervous system [96].

Structural MRI is today widely used for the in-vivo profiling of potential drug candidates for Alzheimer's disease. This was enabled by initiatives such as the Alzheimer's Disease Neuroimaging Initiative, which organised a multi-site program to develop the standards of image acquisition, quality assurance and analysis of this imaging modality that was then approved by regulators [172]. Like EEG, however, standards for the usage of fMRI in drug development have not yet been established, and fMRI studies conducted by pharmaceutical companies have so far only aimed to validate the imaging technique using well-characterised licensed compounds. To avoid large capital investment in technology with an uncertain future, many pharmaceutical companies are choosing to partner with academics

with access to MR scanners rather than install a scanner in-house in order to conduct these studies [196]. An example is Astra Zeneca, who collaborated with the University of Manchester and University of Oxford in 2010-2011 to compare the effect of compound AZD6765 on depressive subjects using fMRI [7].

GSK are one of the few pharmaceutical companies that run fMRI experiments in-house: with financial backing from Imperial College, London and the Medical Research Council, they founded the Clinical Imaging Centre at Hammersmith Hospital, London in 2007. In addition to the £46 million that they invested to build the Centre, they have promised to spend £11 million a year on the development and application of imaging for drug development [115].

7.3.3 Market requirements

For both the clinical and drug development markets, the choice of imaging modality is a balance of cost and quality.

The costs of healthcare in Western society are increasing at an unsustainable rate: as an example of this, $0.46 of every tax dollar was spent on healthcare in Canada in 2012, and this has been predicted to rise to $0.70 by 2022 [70]. For this reason, it is important that new therapies can demonstrate that the savings they bring over the patient's lifetime are greater than their costs.

In addition to the capital and direct running costs, it is important to also consider the following costs of an imaging solution for diagnosing and monitoring neurodegenerative patients:

- Staff time - Solutions that can be used at the hospital bedside do not take up staff time transferring patients to another location in the hospital and off their beds, which can be very time-consuming if the patient has limited mobility. Those solutions that are fast to set up and run will also take up less staff time.

- Staffing - Costs will be minimised if the solution is simple enough that it can be operated by staff without dedicated qualifications to use it.
- Repair - Simpler solutions can be repaired quickly by hospital technicians, rather than needing to be repaired by the manufacturer.

These costs should be weighed up against the savings that a solution for diagnosing and monitoring neurodegenerative patients may enable:

- Patients are able to to remain in employment for longer because earlier diagnosis allows therapies that delay disease progression to be prescribed.
- Patients are able to live more independent lives because more accurate diagnosis means prescribed therapies are more likely to be effective.
- The costs resulting from ineffective treatments being prescribed are reduced, as are the number of clinical visits resulting from diagnoses not being correct or sufficiently personalised.

The ability of an imaging solution to make earlier and better diagnoses depends on how information-rich its signals are, and this measure of signal quality is a function of spatial localisation, temporal resolution and sensitivity. The impacts that earlier and more accurate diagnosis (enabled by greater signal quality) will have on patient quality of life are obvious, however in the clinical setting there are additional measures of quality that need to be considered, such as the solution's portability and sensitivity to motion, because these factors affect the patient experience.

When considering a tool for investigating neurovascular coupling in drug development, the forces of cost and quality are again very interdependent: the quality of the imaging solution in terms of information content will, to some extent, determine the accuracy of the GO/NO-GO decisions made using the data recorded from it, and this will be a strong indicator of profit. However, there are

no extra measures of quality to consider, and so the balance to be considered by pharmaceutical companies is more straightforward than in the clinical setting: will the increase in drug sales resulting from adoption of this new solution be greater than its capital and direct running costs? Given that PET already has an established role in drug development for offering cellular-level information about novel drugs, the greatest value would be added by supplementing this information with systems-level information.

7.4 Recommendations

The incentives for using more information-rich imaging modalities (those with higher spatial localisation, temporal resolution and sensitivity) to investigate neurovascular coupling are greater for pharmaceutical companies than medical centres. The main reason for this is that for new imaging techniques to be adopted by either market, they must demonstrate that they improve diagnosis/drug development beyond current levels. Methods such as PET imaging are already well-established tools for drug development, whereas in the clinic, diagnosis of neurodegenerative disease is made based on interviews, and therefore the information content of any new imaging solution must currently be greater for drug development than clinical use. In addition, the direct rewards that an imaging modality brings to drug development are strongly correlated to its quality (more efficacious drugs are likely to generate more revenue), whereas in the clinic, many of the rewards of using more information-rich imaging are only indirectly felt by medical centres as it is the patient and society that feel the direct benefit (earlier and improved diagnoses improve patient prognosis, meaning they are able to be more productive to society for longer). Therefore, because the information-rich modalities are more costly to acquire and run, the direct cost-benefit ratio is greater for clinics than pharmaceutical companies.

These differences in cost-benefit mean that the optimal hybrid neuroimaging

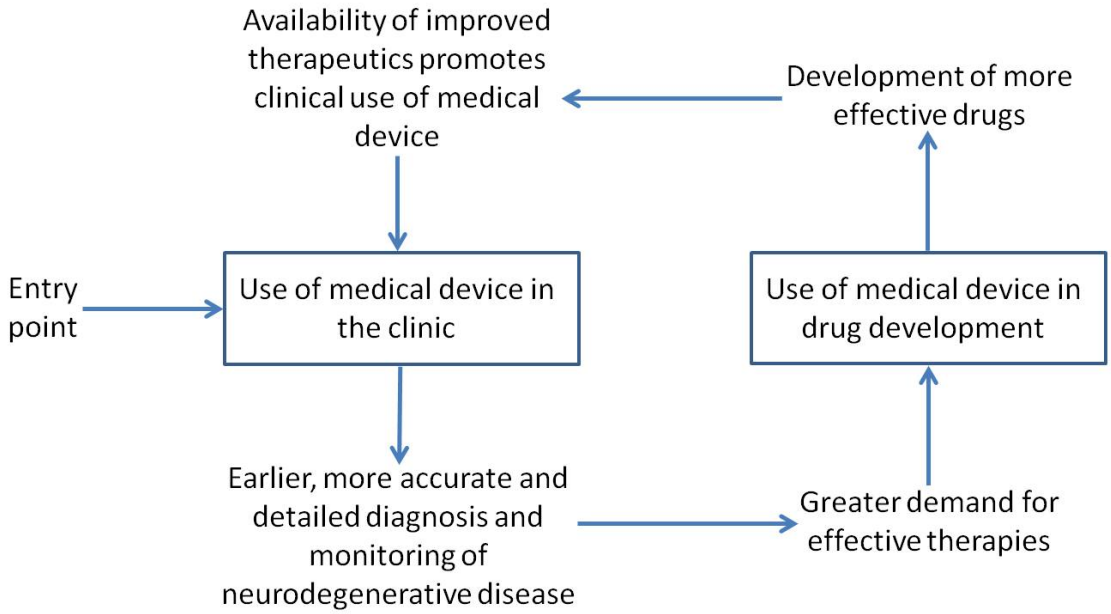

Figure 7.4: The feedback existing between the clinical and pharmaceutical markets means that it is necessary to target both markets. This is because uptake of the software by the clinical healthcare market will only be maximised if effective drugs are available to treat those patients diagnosed and monitored using the software. Equally, uptake of the software by the pharmaceutical industry will only be maximised if patients who would benefit from a new drug are being diagnosed.

solution for the clinical healthcare is different from that for the drug development market. However, because these markets are linked, it is necessary to target both markets, as shown in Fig.7.4.

The recommended route for achieving commercialisation of the generative embedding software is therefore to first develop it for the NIRS-EEG clinical market, enabling people with neurodegenerative diseases to be diagnosed earlier and more accurately, and then to adapt it for entering the pharmaceutical fMRI-EEG market, so that more effective treatments for these patients can be made available.

In addition to this order of initial market entry creating a feedback cycle for the software as soon as both markets are entered, it also makes most sense when one considers where the NIRS-EEG and fMRI-EEG technologies are today. Dartmouth Brain Imaging Group's NIRS-EEG system is a single unit that is undergoing

active clinical testing and which is already contracted to be sold as part of Rogue Research Inc.'s BrainSight family of neuronavigation products, whereas fMRI-EEG requires the integration of two separate products, neither of which is experiencing routine use in drug development yet. It is therefore possible to exploit an already established route to market for the NIRS-EEG version of the generative embedding software, whereas the conditions for releasing the fMRI-EEG version are likely to be more favourable in a few years' time.

Chapter 8

Conclusions

The goal of this thesis has been to:

"describe, demonstrate and evaluate the novel use of a mechanistic method based on generative embedding for making comprehensive, relevant and interpretable conclusions about the implications of neurovascular coupling modulations for BOLD imaging".

This goal has been met by making a number of novel contributions to the field of BOLD-fMRI imaging. The first novel contribution was made in Chapters 2 and 3, where a method for precisely quantifying the extent to which standard BOLD GLM analyses are confounded by the effects of neuromodulations on neurovascular coupling that does not require additional experimental data to be collected was described and demonstrated. This method is based on generative embedding incorporating Combined DCM[1] - the first biophysical model of neuronal and hemodynamics processes in which the correlations between the multimodal processes are estimated.

These two novel contributions were demonstrated using anaesthesia as a case study in neuromodulations. Apart from Alpha-Chloralose and Fentanyl, the effects of all the anaesthetics on neurovascular coupling caused the neuronal activity to be detected with inaccurate statistical significance: neuronal activity levels were

[1]The potential benefits of Combined DCM are not limited to investigating the implications of neuromodulations for neuroimaging, e.g. the Variational Bayes framework used to fit Combined DCM could be used to further our understanding of the specific processes involved in neurovascular coupling.

underestimated under Pentobarbital and Propofol, but overestimated under Ketamine and Isoflurane. These results suggest that if the aim of a BOLD study were to investigate the effects of anaesthesia on neuronal activity, i.e. if anaesthesia were the interest neuromodulation[1], the conclusions from standard BOLD GLM analyses would likely only be accurate for the Alpha-Chloralose and Fentanyl groups (and more likely to be accurate for the Fentanyl group than for the Alpha-Chloralose group). In order for the conclusions from other anaesthetic groups be accurate, the BOLD GLM would need to be modified to compensate for their effects on 1) enhancement of nonlinearities, and 2) hemodynamic impulse response shape.

The aim of Chapter 4 was to understand whether anaesthetic modulation of hemodynamic nonlinearity and/or impulse response shape was responsible for neuronal activity being detected with inaccurate statistical significances in Chapter 2. Using generative modelling, it was found that anaesthetic modulation of impulse response shape, rather than anaesthetic enhancement of hemodynamic nonlinearities, explains most of the mismatch between GLMs and the real responses. Generative embedding of the Combined DCM posterior parameters in linear regressions was then used to identify which specific biological actions of the anaesthetics caused the mismatches of greatest consequence for the GLM analyses and, surprisingly, the biological actions identified were mostly related to neuronal processes, i.e. they impacted on hemodynamics processes indirectly.

Although the combined interest-nuisance neuromodulatory space that may confound BOLD GLMs spans all possible pathologies, ranging from brain diseases such as stroke [42], alterations in neurotransmitter systems such as serotonin in psychiatric diseases [111], and systemic diseases such as diabetes and hypertension [33], to pharmacological substances, such as aspirin [10], anti-depressants [127] and coffee [121], and factors such as age [43], it is usually possible to control most of

[1]The reader is reminded that when nuisance neuromodulations are present, the confounding effects of their interaction with the interest neuromodulation should be considered instead.

them, for example by only including subjects from a small age range or by advising human BOLD-fMRI patients not to consume caffeine before their scans (unless of course the interest neuromodulation of the study is age or caffeine). Generative embedding has the potential to identify which biological processes are affected by these commonly occurring neuromodulations to the extent that BOLD GLM Z-values are impacted upon, in order to create neuromodulation-specific constraints that augment the otherwise less effective and/or computationally expensive correction algorithms that have already been developed. This is the third novel contribution made to the field: a method informing how best to modify the standard GLM to compensate for the confounding effects of neuromodulations without complicating the standard (ANOVA) way of reporting its results.

Instead of investigating the problems that neuromodulations present for predicting neural activations from GLMs because of their effects on neurovascular coupling, the overarching aim of Chapters 5 - 7 was to investigate methods for exploiting the actions of disease neuromodulations on neurovascular coupling to improve disease diagnosis, monitoring and understanding. Using the interesting case study of neurodegenerative disease, it was shown in Chapter 5 that there are significant differences in the linear couplings of some of the multimodal shape features extracted from measures of electrical activity and hemodynamics of these brains compared to healthy brains. This indicates that the prototype hybrid neuroimaging system already developed by Prof. Diamond's laboratory has some power to diagnose neurodegenerative-diseased patients, as it is supported by software that uses the traditional analysis method only (linear regression of the multimodal shape features).

Chapter 6 describes and demonstrates the fourth novel contribution made in this thesis to the field of neuroimaging: the first attempts to understand the limits of hybrid neuroimaging hardware, such as that developed in Prof. Diamond's laboratory, for providing information about neurodegenerative disease state.

Extending the k-nearest neighbours algorithm in order to calculate the mutual information and synergy of spaces involving categorical variables with extremely low bias, it was found that 1) neurodegenerative-diseased and healthy multimodal DCM parameter spaces are most discriminable for late-stage neurodegenerative disease, and 2) the addition of neurovascular coupling information improves the discriminability of early-stage neurodegenerative-diseased brains from healthy brains more than at the later stages. The latter finding supports the development of hybrid imaging systems for early-stage neurodegenerative disease diagnosis, monitoring and understanding, when their value is greatest because of the greater susceptibility of disease progression to early interventions. The commercialisability of these systems to both neurological clinical care and pharmaceutical research was then researched in Chapter 7 and it was discovered that the market forces for a NIRS-EEG system are greater for the clinical market, whereas the market forces for a fMRI-EEG system are greater for the drug development market.

8.1 Evaluation of Chapters 2 - 4

It is important to remember that the insights gained from Chapters 2 - 4 are specific to the anaesthetics dataset collected by Franceschini et al., and therefore they should not be viewed as applicable to all preclinical data collected under these anaesthetics. This is because anaesthetic agent is just one of many neuromodulatory factors, including dosage, administration route, time since administration and stimulation paradigm. What is more, the impacts of these neuromodulatory factors will be both brain-region and species-dependent. Therefore, this work serves only to demonstrate the need to carry out comprehensive and relevant investigations to ascertain the potential of any new experimental modulation to confound GLM analyses by altering neurovascular coupling.

To highlight the specificity of the results from Chapters 2 - 4 to Franceschini et

al.'s anaesthetics dataset, a literature search of the effects of the six anaesthetics on the processes modelled by Combined DCM (neuronal activity, metabolism and hemodynamics) was conducted. As can be seen from Table 8.1, conflictory conclusions about the effect of an anaesthetic on these process have been reported in the literature, resulting from different anaesthetic doses, administration routes and animal models having been used.

Anaesthetic/ Drug	**Neuronal processing**	**Metabolism**	**Hemodynamics**
Alpha-Chloralose	U-shaped potentiation of GABA-receptor mediated currents [124][193] (large IPSP intensity and decay time [170]) caused by increased affinity for GABA, enhanced GABA-induced chloride conductance and increased GABA efficacy [65]. [170] also observed depression of excitatory potentials (8.3μM, guinea-pig olfactory cortex cells), but [193] did not (1 to 100μM, rat cortex cells).	[47] observed no change in cerebral glucose metabolism (60 to 120mg/kg IP, rats), whereas [124] observed a decrease (50mg/kg IV plus 40mg/kg/hr infusion, rat) and [189] observed an increase (80mg/kg IP, rat).	Reduced baseline CBF [124] and CO_2 reactivity [166], which restored to normal levels after a few hours [14].
Pentobarbital	U-shaped potentiation of GABA-receptor mediated currents (large IPSP intensity and long decay [170]), caused by lengthened open lifetime of channels [148]. Also depression of excitatory potentials [170] from reduced glutamate release [39].	Decrease in glucose metabolism found to be significant by [141] (50mg/kg IP, rats), but not [38] (15mg/kg IV, rats) and [189] (40mg/kg, rats). Insignificant decrease in $CMRO_2$ [75].	Decrease in baseline CBF, which [137] (40mg/kg IP, rats) and [140] (50mg/kg IP, [75] rats) found to be significant, but (50mg/kg/hr IV, rats) did not.
Ketamine	[201] (<3μM, Xenopus cells) found that GABA-induced currents were unchanged, whereas [102] (2.5 to 10mg/kg IV, cat) and [170] (0.3mM, guinea-pig olfactory cortex cells) observed small increases resulting from lengthening of IPSP decay rate, and [99] observed significant potentiation of the currents. [170] also found that excitation was reduced, caused by blockage of glutamate receptors [102].	Decrease in glucose metabolism in somatosensory cortex, which [38] (10 to 30mg/kg IV, rat) found to be significant, but [30] (100mg/kg IP, rat) did not.	[30] (100mg/kg IP, rats) and [135] (1mg/kg IV, rabbits) observed significant increases in baseline CBF, but [95] (50mg/kg IP, rats) did not. CO_2 reactivity is also reduced [123].
Fentanyl	Reduced GABA-receptor mediated currents, caused by reduced GABA release, IPSP intensity [89] and function of GABA receptors [200]. [200] also found that EPSP amplitudes were increased.	Small, sigmoidal decrease in $CMRO_2$ [75][29].	Very small decrease in baseline CBF [75] [29].

Isoflurane	[76] (0.064 to 0.096mM gas, rat) found that GABA receptors were potentiated, resulting in lengthening of IPSP decay, whereas [167] (gas, Drosophila) and [12] (1.5 to 3% gas, rat hippocampus cells) report reduced glutamate release resulting from decreased presynaptic excitability.	Significant decrease in glucose metabolism [106][136] and $CMRO_2$ [191].	[106] (0.7 to 2.8%, rats) found that baseline CBF was maintained, whereas [176] (2% gas, rats), [94] (0.7 to 2.8% gas, rats) and [75] (1% gas, rats) found that it was increased, and [191] (baboons) found that decreased below 0.5%. [46] (1.6% gas, cats) and [94] (0.7 to 2.8%, rats) found that CO_2 reactivity was maintained, whereas [176] (2% gas, rats) found that it was reduced. Vasodilation linked to elevated NO levels [9].
Propofol	Potentiation of GABA-receptor mediated currents [168], caused by lengthening of channel open time [71].	Decreased $CMRO_2$ [138][5][183].	Decreased baseline CBF [138][5][183], but whereas [5] (12 to 48mg/kg/hr infusion, dogs) and [183] (0.2mg/kg/min infusion, humans) found that CO_2 reactivity was maintained, [50] (2.5mg/kg bolus plus 150mg/kg/min infusion, humans) found that was reduced. [183] also reported an increase in vascular resistance.

Table 8.1: Neuromodulatory effects of the anaesthetics in the anaesthetics dataset, informed by a literature review.

What is more, most of the findings listed in Table 8.1 were not gained from evoked responses, but it is important to also consider the interaction of anaesthetic and stimulation paradigm for this work as Combined DCM requires perturbation from evocation. This interaction is complex because different stimulation frequencies suppress and/or potentiate different neural cell populations and alter hemodynamic nonlinearities, just like different anaesthetic modulations have been found to do. Therefore, the conclusions drawn from Chapters 2, 3, and 4 are also specific to the stimulation paradigm used to evoke the responses that Combined DCM was fitted to (3Hz trains of 200μs, 1.3mA stimuli lasting 5s).

8.1.1 Future work

The generative embedding framework proposed in this thesis for understanding whether a modulation alters neurovascular coupling to the extent that GLMs detect neuronal activity with inaccurate significance levels provides greater insight than any other method known to the author. This is because it tests hypotheses about GLM significances, which are the statistics reported in psychology and cognitive neuroscience BOLD research, rather than hypotheses about less relevant statistics such as correlation coefficients.

However, it is appreciated that the integrity and suitability of Combined DCM is the greatest weakness of this framework in its current form. The author proposes that the relatively simple changes detailed in Appendix E be made to Combined DCM to overcome some of these weaknesses in her future work.

One weakness that would be considerably less simple to overcome, however, is that the generative embedding procedure used in this thesis only considers the effects of modulations on the temporal dynamics of neurovascular coupling, yet the hemodynamic response is a spatiotemporal convolution of neuronal activity and it is likely that the spatial coordination of neural biology is also susceptible. Spatial modulation of neurovascular coupling may have implications for 4D hemodynamics datasets analysed using GLM, causing differences in the spatial distributions of

brain activations calculated for subjects with otherwise identical brain activation patterns. Ways in which neuromodulations may impact on the spatial pattern of brain activations include [112]:

1. How capillaries detect focal excitatory and surrounding inhibitory neuronal activity,
2. How vasoactive signals are transported to upstream arteries,
3. How CBF between active and inactive brain regions are balanced.

For example, Shtoyerman et al., 2000, found that although CBV was larger in unanesthetised animals, the increases were more localised, and this could have been caused by reduced diffusion of vasodilatory signals or enhanced surround inhibition [174].

Using the generative embedding algorithms presented in this thesis to explore the effects of neuromodulations on the spatial pattern of neural activations estimated using GLM would increase the computational expense of Combined DCM fitting enormously. This is because it would be necessary to balance the fitting of Combined DCM to each voxel of data with the fitting of a spatial model that enables processes such as diffusion of vasoactive substances across neighbouring voxels to occur.

8.2 Evaluation of Chapters 5 - 6

Administration of LPS induces a well-established model of neuroinflammation with published effects on brain function [36]. In addition, because this model of neuroinflammation is simple to induce, especially given that it is a model routinely used in the Sibson laboratory group to which I am affiliated, and has many fewer animal welfare issues than more disease-specific transgenic models, it was deemed the most suitable model for the proof-of-concept work reported in Chapters 5 and 6.

Unfortunately, the LPS model did not induce the specific brain disease state of early-stage neurodegenerative disease that is of greatest interest to the author because of the unique neurovascular uncoupling that is hypothesised to occur (modulated hemodynamics but relatively unmodulated neuronal activity). What is more, LPS causes acute inflammation whereas neurodegenerative disease is a chronic condition, meaning that the specific profile of the neuroinflammation may not be wholly relevant to neurodegenerative diseases.

An important next stage to this work would therefore be to identify a more optimal animal model of neuroinflammation. This will undoubtedly require a balance to be struck between logistical and experimental complexity and model relevance.

8.3 Evaluation of Chapter 7

In Chapter 7 the commercialisability of the generative embedding algorithms developed in this thesis was considered. The two obvious markets for software implementing these algorithms were researched:

1. Clinical, for improved diagnosis and monitoring of neurodegenerative diseases,
2. Pharmaceutical, as a tool for the development of drugs for neurodegenerative diseases.

After analysing the requirements of these two markets in conjunction with the properties of the individual neuroimaging modalities available today, it was recommended that a NIRS-EEG version of the software be developed first, to sell into the clinical healthcare market, followed by a fMRI-EEG version, to sell into the drug development market. By selling to both markets, positive feedback should occur as earlier and more accurate diagnosis and monitoring of neurodegenerative patients increases the market for effective drugs to treat them,

Property	NIRS-EEG	fMRI-EEG	MEG-Susceptometry
Spatial localisation	Low	Very high	High
Temporal resolution	Very High	Medium	High
Capital and running costs	Low	High	Very high
Portability	High	None	None
Motion artifacts?	No	Yes	Yes
Simplicity	High	Low	Low

Table 8.2: Comparison of NIRS-EEG, fMRI-EEG and MEG-susceptometry (predicted). The information available in the signals is a function of the spatial localisation and temporal resolution. Modalities sensitive to motion artifacts require the patient to be very still during imaging, which is not possible for some patient populations. Simplicity refers to simplicity to operate and repair the hardware (and any associated software).

and having effective drugs available will encourage physicians to diagnose and monitor their patients earlier and more accurately.

8.3.1 Neurovascular imaging modalities of the future

A number of imaging techniques may prove to be competitors to NIRS-EEG and/or fMRI-EEG, although most are very under-developed at this time. One example is imaging both neuronal activity and hemodynamics using NIRS, as NIRS has been shown to be sensitive to membrane potential (neuronal activity) as well as blood oxygen concentration (hemodynamics) [159]. A second example is MEG-Susceptometry, which is also being led by Dartmouth Brain Imaging Group. MEG-Susceptometry combines the superconducting quantum interference devices used to measure the magnetic field resulting from neuronal activity (MEG), with ultra-low field electromagnets to measure the paramagnetic deoxyhemoglobin content (Susceptometry) [22]. The pros and cons of NIRS-EEG, fMRI-EEG and MEG-Susceptometry are listed in Table 8.2.

In order to remain a competitive software provider to both clinics and pharmaceutical companies in the future, it will be necessary to monitor the neuroimaging landscape for these and other emerging modalities and quickly respond with new versions of the generative embedding software for any that have

the potential to be more attractive to the clinical and drug development markets than NIRS-EEG and fMRI-EEG, respectively.

8.4 Summary

It is widely recognised that in order to gain comprehensive and accurate understanding of biology, sophisticated analysis techniques need to be developed. In this thesis, the novel use of a mechanistic method based on generative embedding for making comprehensive, relevant and interpretable conclusions about the implications of neurovascular coupling modulations for BOLD-fMRI imaging has been demonstrated. When used to quantify the extent to which interest and/or nuisance neuromodulations confound the standard BOLD GLM framework, generative embedding identified neuromodulations that cause neuronal activity to be detected with inaccurate statistical significances, which the commonly used traditional analysis method could not detect [56]. Using generative embedding once again, the large range of neuromodulatory biological actions causing these confounding effects were explicitly identified, highlighting the need for generative modelling and/or generative embedding to be used to improve the effectiveness and practicality of existing methods for compensating for neurovascular coupling modulations in the standard BOLD GLM.

Stanimirovic and Friedman, 2012, identify "a key future challenge in multiparametric disease evaluation and staging as being meaningful integration of disease-specific cerebrovascular biomarker profiles with functional (neuro)vascular assessments using advanced imaging techniques" [182]. This thesis has gone some way to meeting this challenge: the potential of hybrid neuroimaging systems supported by generative embedding algorithms for the diagnosis, monitoring and understanding of brain disorders and diseases accompanied by neurovascular uncoupling has been investigated. The value of hybrid neuroimaging systems over established, separate neuroimaging technologies, has been found to be greatest for

early-stage neurodegenerative disease, and strategies to commercialising this technology have been developed.

Appendix A

BOLD GLM

The canonical HRF is a function used to describe the temporal dynamics of oxygenated blood flux to a single, short stimulus. In SPM, which is optimised for human BOLD imaging, it is generated from the double-gamma variate, which is shown in Fig.A.1(b) along with its temporal and dispersion derivatives. In this thesis, GLMs are fitted to the [ΔHbR] responses generated by awake rats and so the canonical [ΔHbR] and its temporal and dispersion derivatives shown in Fig. A.1(b) are the basis functions used instead.

Using these basis functions to create explanatory variables, the resulting linear model for the de-meaned BOLD signal in one voxel collected from subject i, y_i, is:

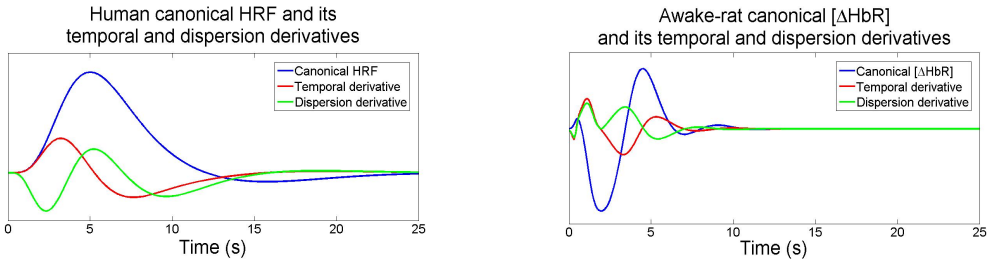

(a) (b)

Figure A.1: Canonical [ΔHbR] and its temporal and dispersion derivatives (a) commonly used as basis functions in BOLD GLMs in standard human BOLD imaging, and (b) used in this thesis, as [ΔHbR] responses from rats are analysed instead.

$$y_i = \sum_{j=1}^{n} \beta_{ij} x_{ij} + \epsilon_i, \tag{A.1}$$

where β_{ij} = regression coefficient for explanatory variable x_{ij} (same dimensions as y_i), which models the BOLD response to stimulation type j, and $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \sigma^2 \boldsymbol{I})$, which aims to capture the data unexplained by the j other explanatory variables.

The regression coefficients are estimated using matrix manipulation so that a least-squares solution is calculated:

$$\hat{\boldsymbol{\beta}}_{\boldsymbol{i}} = (\mathbf{X_i^T X_i})^{-1} \mathbf{X_i^T} y_i, \tag{A.2}$$

where $\hat{\boldsymbol{\beta}}_{\boldsymbol{i}}$ = estimated regression coefficients and $\boldsymbol{X_i}$ = design matrix containing all j explanatory variables for subject i.

Because the anaesthetics dataset has [ΔHbR] recorded from a single location only, the GLM analyses carried out in Chapters 2 and 3 are equivalent to that carried out if only a single voxel of BOLD data had been collected, removing any need to filter the data spatially, apply motion correction[1] (because it is not possible to calculate the motion of a single datapoint) or account for multiple comparisons. It was, however, necessary to pre-process the data in two other ways: 1) highpass filtering[2] in order to remove the slow-frequency variations (drifts) unrelated to the evoked neuronal responses, and 2) pre-whitening in order to remove the autocorrelations[3].

The next step of a neuroimaging GLM analysis is to answer questions about the significances of the signal changes that the GLMs have been fitted to. These

[1]The different motion correction algorithms implemented in the different major software have been shown to at least partly account for GLM sensitivity, with the algorithm in SPM giving significantly higher sensitivity than FSL, ANFI or Brainvoyager [119].

[2]In SPM the cut-off of the default highpass filter is 128Hz, but the default filter threshold varies amongst the major software.

[3]It is important to prewhiten the data to ensure the least-squares solution to Equ.A.2 is the maximum likelihood estimate, as this only occurs when $\boldsymbol{\epsilon}$ are independent and identically distributed. In SPM, the autocovariance matrix is estimated by fitting an autoregression model using restricted maximum likelihood. In the other major fMRI software packages Tukey tapering or pseudogeneralised least squares are used instead [199] [21]

ANOVA calculations are carried out by creating contrasts (linear combinations), $\boldsymbol{c}$, of the regression coefficients that test questions about relative activations. The signifiances of these relative activations are calculated using:

$$\mathbf{t} = \frac{\mathbf{c}^T\boldsymbol{\beta}}{\sqrt{Var(\boldsymbol{\epsilon})\mathbf{c}^T(\mathbf{X}^T\mathbf{X})^{-1}\mathbf{c}}}. \tag{A.3}$$

This t-value, which is later converted to a Z-value in the major software packages to enable multi-subject comparisons (in a standard space), is determined by the magnitude of the canonical [ΔHbR] $\boldsymbol{\beta}$ (effect size, the numerator) and the standard error of the resulting GLM (the denominator). Therefore, the contrast when only canonical [ΔHbR] is used as a basis function in the response explanatory variable is [1], whereas it is [1 0 0] when the temporal and dispersion derivatives are also included[1].

In the SPM software, the second-level (group-level) Z-values are a function of the subject-specific effect sizes from the first-level analyses (individual level) only, i.e. they are not sensitive to the intra-subject variances or to the first-level standard errors[2].

[1] It is important that the $\boldsymbol{\beta}$ for the temporal and dispersion derivatives are not included in the contrast because the aim of the derivatives is to regress out the variances captured by the derivatives.

[2] In order to fully take the response variability across rats into account, the second-level statistics should be a function of the intra-subject variances as well as subject-specific effect sizes, and this is the case for some fMRI software packages. The authors of SPM decided against incorporating the intra-subject variances into second-level analyses because of the increase in algorithmic complexity that it would involve.

Appendix B

Fitting DCMs

DCMs are fitted to observational data using Variational Bayes. The first step of this algorithm is to generate a predicted response using the prior means as starting values for the parameters. These prior values are guesses, informed by experiments and/or analytical observations from previous modelling, as to what these parameter values are.

B.1 State variable integration

In DCMs, the predicted response is a function of the state variables (x_1:x_9 in DCM-EEG and s, f, ν and q in mDCM-fMRI, see Fig.3.2 and 3.3, respectively), and the exact dynamics of these state variables are specified by ordinary differential equations (of the form $\dot{x}(t) = f(x, u)$). For a given set of parameter values, it is possible to estimate these nonlinear dynamics at each timestep:

$$\tilde{X}(t_{k+1}) \approx e^{J\tau_k}\tilde{X}(t_k), \tag{B.1}$$

where $J = \frac{\partial f}{\partial x}$ (the system's Jacobian) and $\tau_k = t_{k+1} - t_k$ is the integration timestep [41].

Because of the differing timescales of neuronal activity and hemodynamics, the

algorithms used to make these estimations are different for DCM-EEG and DCM-fMRI.

For DCM-EEG, it is important to preserve the nonlinearities in the state variables because their dynamics are very fast. Therefore, to speed up integration of the state variables without ignoring their nonlinearities, J is approximated by its value at the fixed point, X_0 (when the state variables are unperturbed), meaning that it only has to be calculated when new exogenous input, u, arrives, i.e. τ_k are time periods over which no new exogenous input arrives. As stimulus functions are normally sparse, this quasi-analytical integration scheme can be an order of magnitude faster than straightforward numerical integration.

For DCM-fMRI, integration of the state variables can be sped up further using bilinear approximations to the state variables, X, about X_0, as these are sufficient for capturing the slower temporal dynamics of blood [1]:

$$\dot{X}(t) \approx f(X_0,0)+\frac{\partial f(X_0,0)}{\partial X}(X-X_0)+\sum_i u(t)_i \left(\frac{\partial^2 f(X_0,0)}{\partial X \partial u_i}(X-X_0)+\frac{\partial f(X_0,0)}{\partial u_i}\right). \tag{B.2}$$

B.2 Variational Bayes algorithm

In both DCM-EEG and DCM-fMRI, the output, $h(u,\boldsymbol{\theta})$, is used to generate a predicted response, $y_{predicted} = h(u,\boldsymbol{\theta}) + N$, which is compared to the true signal, y, being fitted to in order to calculate a prediction error. The likelihood that y was generated by the prior parameter values that gave rise to the predicted response is then:

$$p(y|\boldsymbol{\theta},\lambda) = N(h(u,\boldsymbol{\theta}),\Sigma(\lambda)), \tag{B.3}$$

where λ is the error covariance hyperparameter (estimated using the same

restricted maximum likelihood algorithm used to estimate the autocorrelations for GLM mentioned in Appendix A). Now that the autocorrelations in the true signal are accounted for, it is assumed that the prediction error is Gaussian only. This assumption forms the basis of the E-step of the Expectation-Maximisation (E-M) algorithm that is then used to repeatedly jointly update $\boldsymbol{\theta}$ and λ so that the predicted response approaches y under model complexity constraints.

The aim of the E-step is to update the parameter moments so that they maximise both the likelihood and prior potential given the current estimates of the hyperparameters. This is made possible by approximating the output by its linear form:

$$h(u, \boldsymbol{\theta}) \approx h(\nu_{\theta|y}) + J(\theta - \nu_{\theta|y}), \tag{B.4}$$

and then using the assumption that the prediction errors are normally distributed to calculate the likelihood:

$$p(y|\boldsymbol{\theta}) \propto e^{\frac{-1}{2}(r-J(\theta-\nu^i_{\boldsymbol{\theta}|y}))^T C_E^{-1}(r-J(\theta-\nu^i_{\boldsymbol{\theta}|y}))}. \tag{B.5}$$

Because Baye's rule is given by:

$$p(\boldsymbol{\theta}|y) \propto p(y|\boldsymbol{\theta})p(\boldsymbol{\theta}), \tag{B.6}$$

where $p(\boldsymbol{\theta})$ = prior probability, it can be seen that multiplying the likelihood with the Gaussian prior probabilities gives the following expression for the posterior parameters, $p(\boldsymbol{\theta}|y, \lambda)$:

$$p(\boldsymbol{\theta}|y) \propto e^{\frac{-1}{2}(\theta-\nu^{i+1}_{\boldsymbol{\theta}|y})^T C^{-1}_{\boldsymbol{\theta}|y}(\theta-\nu^{i+1}_{\boldsymbol{\theta}|y})}. \tag{B.7}$$

What is important is that this expression is recursive, meaning that it can iteratively be used to update the parameters so that they maximise the likelihood

of observing the data whilst being constrained to stay close to their prior mean values (the penalties for discrepancies are dependent on their prior variances, i.e. how confident the prior mean values are) [58]. Because only the first two moments of the posterior densities are estimated in this algorithm, i.e. the parameter space is assumed to be multivariate Gaussian, this E-M algorithm gives a fixed-form Laplace approximation to the true posterior density.

In the M-step, the aim is to update the hyperparameter moments so that they better explain the autocorrelations in the data given the current estimates of the parameters (improving the strength of the assumption that the prediction error is Gaussian). This is achieved by increasing the negative free energy, F, which is a function of the log-likelihood, l. However, because l depends on parameter density $p(\boldsymbol{\theta})$, which is unknown, it is necessary to integrate it out using an approximate Gaussian conditional density, $q(\boldsymbol{\theta})$ (the multivariate Gaussian parameter space estimated in the E-step) [59]. Fisher scoring is then used to direct the hyperparameter space search:

$$\lambda^{i+1} = \lambda^{i} + \langle \frac{\partial^2 F}{\partial \theta^2} \rangle^{-1} \frac{\partial F}{\partial \theta}. \tag{B.8}$$

It is necessary to iterate the E- and M-steps because the covariance of the parameter estimates is itself a function of the error covariance, meaning that the parameters and hyperparameters need to be estimated jointly [58]. This iteration of the E- and M-steps continues until the hyperparameter estimates converge.

Appendix C

Review of Linking Models

Models that capture the biological processes underlying two or more neuroimaging modalities are most commonly developed and optimised for multimodal data fusion, which aims to make the most of the advantages of each of the individual imaging modalities. For example, in EEG-fMRI fusion studies, the aim is to develop algorithms that allow neuronal activations to be detected more precisely by simultaneously exploiting the high temporal resolution of the EEG signal and the high spatial resolution of the fMRI signal. This aim is different from the aim for Combined DCM in this thesis, which is to interrogate the relationships between neuronal and hemodynamic processes occurring in the same brain region. For this reason, the priority for Combined DCM is model accuracy.

DCM-EEG and DCM-fMRI are both experimentally motivated and they have both undergone rigorous analytical testing. It was therefore decided that only models that 1) use variables calculated in DCM-EEG as input, and 2) generate variables that can drive mDCM-fMRI as output, would be considered for Combined DCM. This meant that Riera et al's model, which links the electro-vascular neural mass model to the Balloon Model, was not considered because it would require DCM-EEG to be reconfigured to calculate Nitric Oxide concentration (a product of neuronal metabolism), as that is what drives CBF in their model.

C.1 Buxton et al. and Babajani et al.'s models

The first model that was considered for linking DCM-EEG and mDCM-fMRI was Buxton et al.'s model. In this model, bloodflow is a function of the convolution of the neuronal activity, N, and a delayed gamma-variate impulse response function, h [26]:

$$f(t) = 1 + (f_1 - 1)h(t - \delta_{t_f}) \star N(t), \tag{C.1}$$

where N is the difference between the excitatory and inhibitory currents, h is equivalent to the impulse response functions P(t) used in DCM-EEG to model excitatory and inhibitory mean transmembrane potentials (see Fig.3.2), and f_1 and δ_{t_f} parameterise the magnitude and delay of the bloodflow in response to N.

The problem with this model is that it cannot capture the initial dip - a fast, transient hemodynamic feature reported to sometimes occur in the first couple of seconds after stimulation and which is clearly visible in many of the [ΔHbR] responses in the anaesthetics dataset (see Fig.C.1 for an example response). The initial dip is believed by most to be caused by rapid increase in oxygen metabolism before CBF increases [108][128], however Buxton et al.'s model makes oxygen metabolism a fixed function of CBF (creation of HbR = $\frac{f(1-(1-E0)^{\frac{1}{f}})}{E0}$), meaning that q cannot change before f[1]. This is also true of Babajani et al.'s model, which is similar to Buxton et al.'s model except that it assumes that the sum of all the postsynaptic potentials drive CBF, as opposed to their difference.

[1] An alternative hypothesis for the basis of the initial dip is that it results from fast increases in CBV. This hypothesis is supported by the experiment carried out by Sirotin et al. in which increases in HbT but not HbR were observed in the first couple of seconds after stimulation [177]. Observation of initial dips in [ΔHbR] signals favours the hypothesis of O_2 consumption occurring before CBF increases, as the only way to cause decreases in [ΔHbR] by CBV dynamics would be for CBV to decrease in the first couple of seconds after stimulation, which is contrary to the observation made by [177]

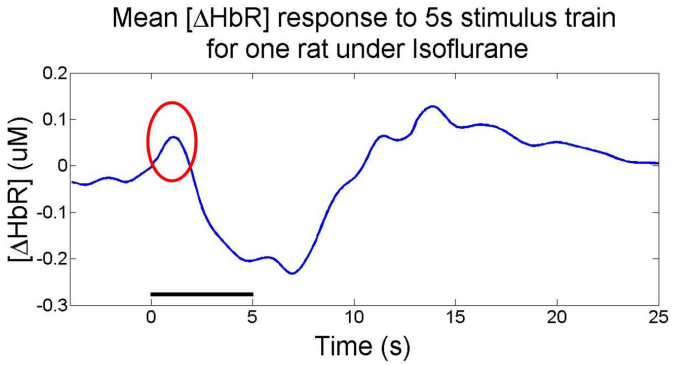

Figure C.1: Average [ΔHbR] response to 5s stimulation for one rat under Isoflurane with the initial dip, which Buxton et al. and Babajani et al.'s models cannot capture, highlighted.

C.2 Sotero et al.'s model

Because initial dips are predominantly visible under only some anaesthetics in the anaesthetics dataset, such as Isoflurane, it is important to capture these transient features as they may contribute to the differences in GLM significances across anaesthetics. Sotero et al.'s metabolism model provides one way to do this as it models oxygen metabolism as a function of the total excitatory synaptic potential, rather than as a function of CBF (full details given in 3.1.2.1). An additional benefit of Sotero et al.'s model over the alternatives tested is that all of the parameters and processes within it have biological interpretations, such as efficiencies and time constants of glucose consumption. This is not the case for Buxton et al. and Babajani et al.'s models, as δ_{t_f} and δ are unspecified delay parameters in these models. This is advantageous because it means that more relevant and specific learning can be made from Sotero et al.'s model when it is part of a generative embedding analysis. For these two reasons, Sotero et al.'s model was chosen to link DCM-EEG and mDCM-fMRI in Combined DCM. The dynamics of this model are shown in Fig.C.2.

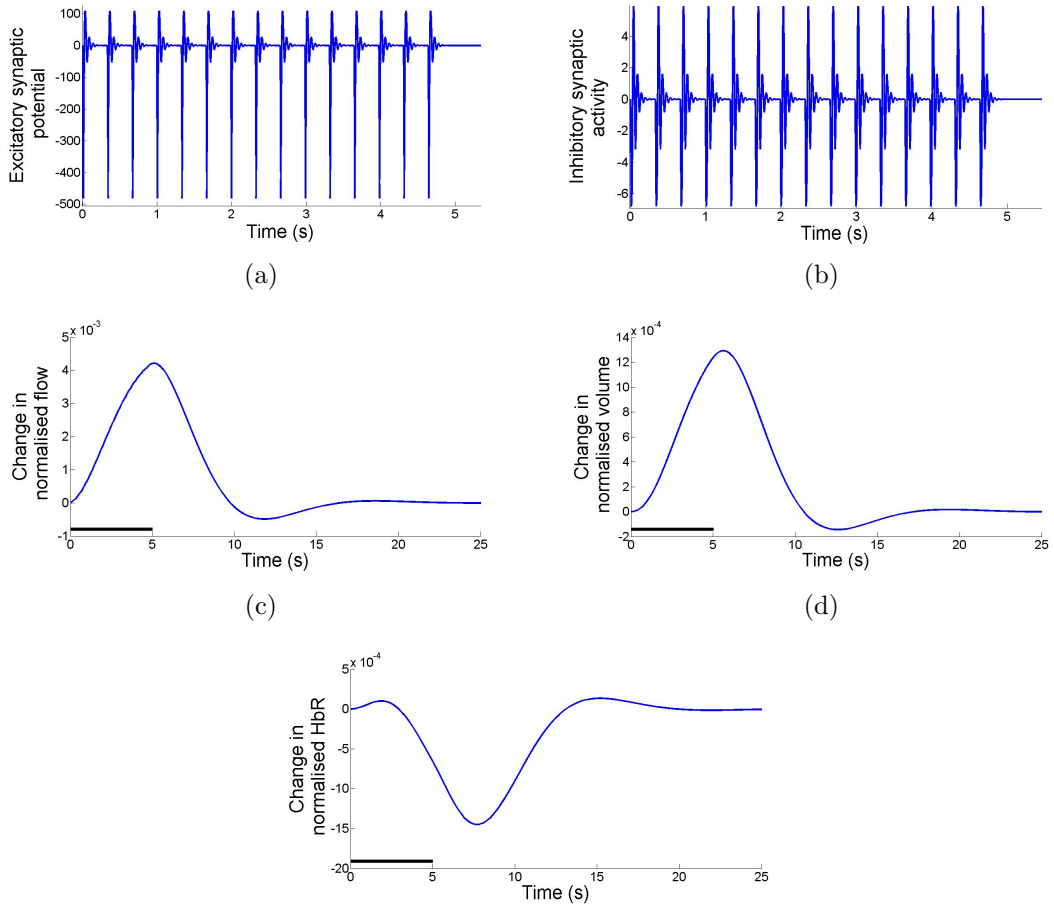

Figure C.2: State variable dynamics in response to a 5s train of stimulation for Combined DCM with Sotero's model linking DCM-EEG and mDCM-fMRI and with parameters set to prior mean values: a) excitatory neuronal activity (zoomed in to see detail in first five seconds, when stimulation occurs); b) inhibitory neuronal activity (zoomed in to see detail in first five seconds, when stimulation occurs); c) change in normalised flow, f; d) change in normalised volume, v; e) change in normalised HbR, q. Notice that even when the parameters are set to prior mean values, a small initial dip is apparent in the change in HbR signal.

Appendix D

Literature review: Bias correction schemes

Panzeri et al. recommend that one of Panzeri-Treves, Quadratic Extrapolation, Best Universal Bound or Nemenmen Shafee-Bialek corrections be combined with shuffling correction to give acceptable variances and unbiased estimates of the density of binned univariate data [145].

Panzeri-Treves and Quadratic Extrapolation corrections both rely on the fact that because the entropy estimated from the data is the maximum likelihood estimate, an analytical correction of $\frac{b-1}{2N}$ removes first-order bias (the leading term in the expansion of $\frac{1}{N}$), where b is the number of bins and N is the number of datapoints (Miller-Maddow correction) [117]. However, this expansion is only true when sampling asymptotically, i.e. for very large datasets, and if none of the bins are empty [188]. Panzeri-Treves uses Bayesian statistics to estimate the true response probabilities in each bin, with a prior that is one constant for all non-zero bins, and another constant for all empty bins. This approach calculates a bias that is sensitive to the fact that more bins may have been occupied had more data been available, whilst allowing for some bins to be 0 if that is considered to reflect the true probability distribution [146]. Quadratic Expansion, on the other hand,

calculates the coefficients of the first two terms of the $\frac{1}{N}$ expansion by comparing the information available in different fractions of the data [188].

Best Universal Bound and Nemenmen Shafee-Bialek do not rely on sampling being asymptotic and therefore are more suitable for smaller data sets. They both reduce bias by generating more uniform bin probabilities (meaning that there are no empty bins, which create particularly high levels of bias). Best Universal Bound achieves this by approximating the binned distribution using binomial polynomials [142]. Nemenmen Shafee-Bialek, however, uses Bayesian statistics in a similar way to the Panzeri-Treves correction in order to calculate P(p|n) - the probability of distribution p proposed by the binned data given the actual data n - but this time weighting the likelihood with a Dirichlet prior, which has a regularising effect [125].

Unfortunately, these bias correction schemes do not correct the biases in estimates of correlations between spaces, meaning that they are insufficient when estimating $H(\mathbf{A}, \mathbf{B})$ and $H(\mathbf{A}, \mathbf{B}|\mathbf{S})$ (entropies of the unconditional and conditional joint spaces, respectively). To overcome this, Montemurro et al., 2007, recommend that the following shuffling correction is also applied when calculating the mutual information of multivariate spaces:

$$I(\mathbf{A}, \mathbf{B}; \mathbf{S}) = H(\mathbf{A}, \mathbf{B}) - H_{ind}(\mathbf{A}, \mathbf{B}|\mathbf{S}) + H_{sh}(\mathbf{A}, \mathbf{B}|\mathbf{S}) - H(\mathbf{A}, \mathbf{B}|\mathbf{S}), \tag{D.1}$$

where $H_{sh}(\mathbf{A}, \mathbf{B}|\mathbf{S})$ is the entropy of the shuffled data, and $H_{ind}(\mathbf{A}, \mathbf{B}|\mathbf{S})$ is the conditional entropy estimated for the marginal probabilities but excluding the contributions of the correlations between A and B [145]. This corrections works because the difference between $H(\mathbf{A}, \mathbf{B}|\mathbf{S})$ and $H_{sh}(\mathbf{A}, \mathbf{B}|\mathbf{S})$ is bias-corrected $H(\mathbf{A}, \mathbf{B}|\mathbf{S})$[1], and the difference between $H_{ind}(\mathbf{A}, \mathbf{B}|\mathbf{S})$ and $H(\mathbf{A}, \mathbf{B})$ is bias-corrected $H(\mathbf{A}, \mathbf{B})$.

Ostwald et al., 2010, combined Panzeri-Treves correction with the shuffling

[1]The bias cancels out because $H_{sh}(\mathbf{A}, \mathbf{B}|\mathbf{S})$ does not include the contributions of correlations between $\mathbf{A}$ and $\mathbf{B}$ to the entropy estimate, but it still possesses similar amounts of bias as $H(\mathbf{A}, \mathbf{B}|\mathbf{S})$ because the same amount of data is used in its estimate.

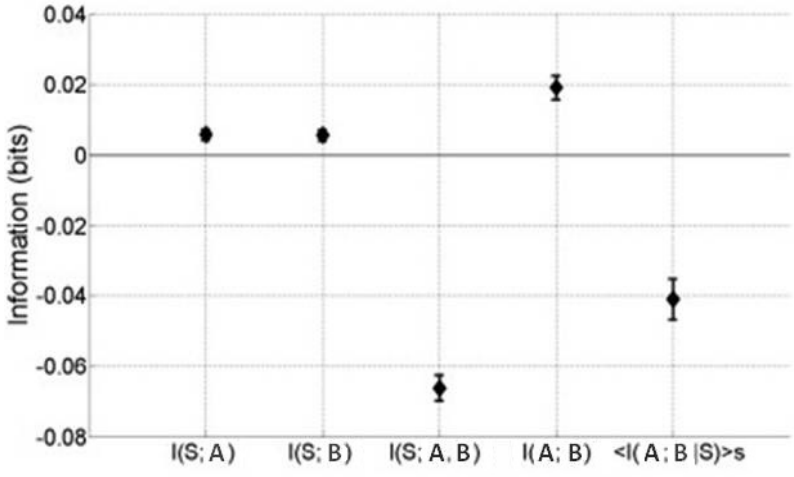

Figure D.1: Estimates of information theoretic values from 100 simulations of model M0 (where the response variables A and B are simulated to be independent of each other and categorical variable **S**) after bias-correction using Panzeri-Treves and shuffling correction schemes. I(**S**;**A**) and I(**S**;**B**) remain close to 0 (as they were before the bias-corrections because they are calculated from the marginal densities only), but the other information theoretic quantities have either not been corrected enough (I(**A**;**B**)), or have been over-corrected (I(**S**;**A**,**B**) and $\langle I(\mathbf{A};\mathbf{B}|\mathbf{S})\rangle$). Figure adapted from Ostwald et al., 2010, Figure 2B [139].

correction but found that they did not satisfactorily correct the systematic errors in the calculations of information theoretic quantities for their M0 model (where their response variables A and B were simulated to be independent of each other and categorical variable **S** - see Fig.D.1 for evidence of this) [139]. Therefore, they further corrected the information values by subtracting the average error from 100 simulations of their M0 model. However, in preliminary work, it was found that when these corrections were subtracted from information theoretic values obtained using the Panzeri-Treves and shuffling corrections, negative information values were often obtained, which are nonsensical. Therefore, the corrections that Ostwald et al. promote are only suitable for data generated by their M0 model. For this reason, binless methods for calculating information theoretic values that are less susceptible to bias were pursued.

Appendix E

Future modifications to Combined DCM

The author intends to implement the following four changes to Combined DCM in order to improve its utility for any future generative embedding work.

E.1 Separating presynaptic excitability and inhibition

Replacement of the sigmoidal function that transforms all presynaptic potentials into firing rates[1], with one sigmoid specific to inhibitory neurotransmitter release and one specific to excitatory neurotransmitter release would allow hypotheses about the effects of neuromodulations on the separate mechanisms of presynaptic excitability and inhibition to be tested. This would result in two extra parameters (S1 and S2 would be replaced by $S1_e$ and $S2_e$, and $S1_i$ and $S2_i$. respectively) in Combined DCM, but the evidence for anaesthetic modulation of presynaptic excitability for at least Isoflurane is compelling [12] and therefore ought to be captured.

[1] A sigmoidal function is often used to transform presynaptic potentials into firing rates because the average of a population of fluctuating neuronal states is a sigmoid [77].

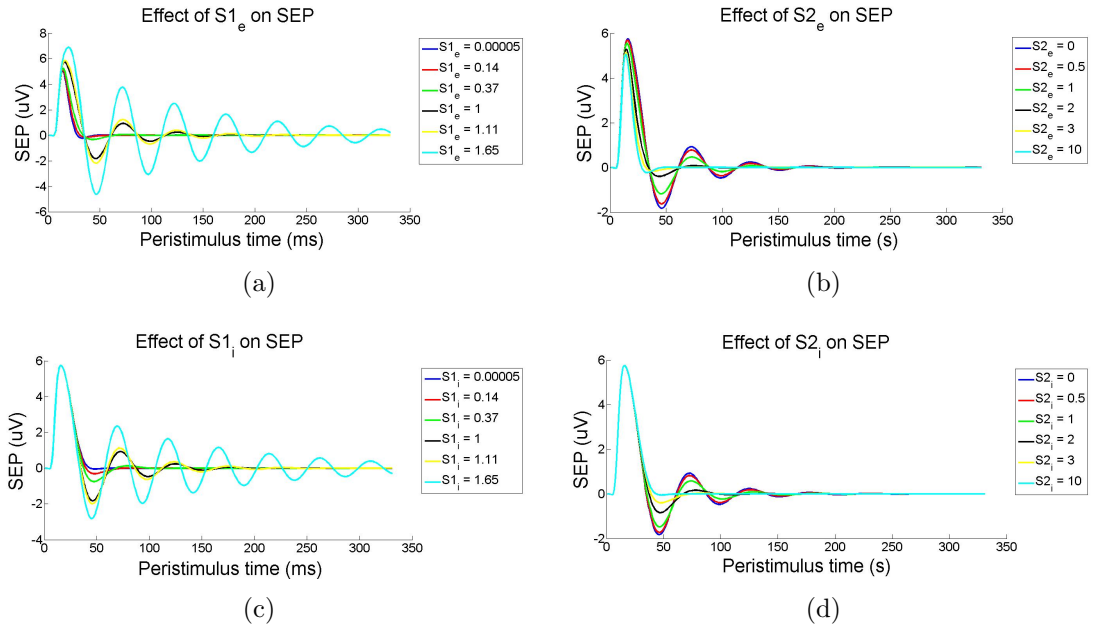

Figure E.1: SEP responses to $S1_e$, $S2_e$, $S1_i$ and $S2_i$ varied between limits that give sensible SEP responses.

The SEP and [ΔHbR] responses for different values of $S1_e$, $S2_e$, $S1_i$ and $S2_i$ are shown in Fig.E.1 and E.2, respectively.

E.2 Capturing neural adaptation

In its present form, Combined DCM fits to the average EEG responses to each train of stimuli. This is because it has no mechanism for capturing variation in the electrical responses to the individual evocations in each stimulation train. A large part of this variation is caused by neural adaptation[1], which results from calcium-dependent potassium channels staying open for longer in response to steady-stimulus current, and therefore longer periods of membrane hyperpolarisation during which time subsequent action potentials cannot occur [51]. The consequence of not modelling adaptation is greater mismatch between the actual EEG responses and the corresponding DCM-EEG fits, and therefore the

[1]Neural adaptation is viewed as the simplest form of learning.

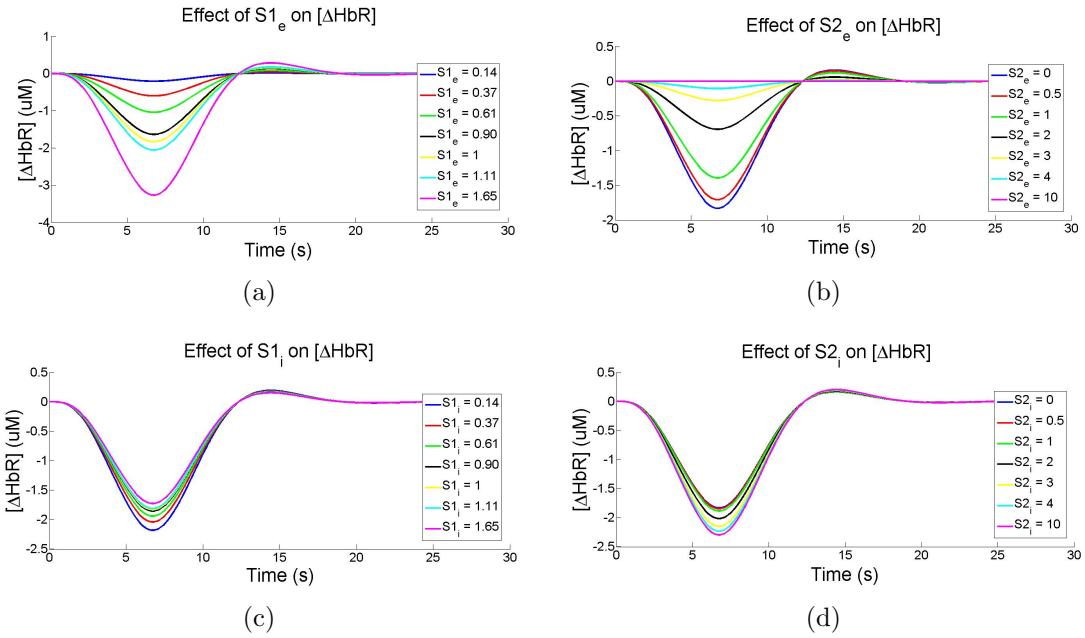

Figure E.2: [ΔHbR] responses to $S1_e$, $S2_e$, $S1_i$ and $S2_i$ varied between limits that give sensible [ΔHbR] responses.

hemodynamics in Combined DCM being driven by incorrect neuronal activity.

As can be seen in Fig.E.3, there are clear differences in neural adaptation under the anaesthetics of the anaesthetics dataset. These patterns for adaptation across anaesthetics match those that Masamoto et al. observed [113], including a prolonged refractory period under Alpha-Chloralose, causing every other response to be supressed.

Laxminarayan et al. simplified and then extended the generative model developed by Riera et al. [161] to incorporate adaptation. Using a different dataset collected by Franceschini et al. [55], in which EEG responses to 4s trains of medial nerve stimuli delivered at 1 to 8Hz were recorded from somatosensory cortex under Alpha-Chloralose anaesthetic, they developed the following set of adaptation rules:

1. Adaptation strength increases as stimuli are delivered at greater frequency:
    - There is little or no adaptation when stimuli are delivered at 2Hz or slower.
    - For stimulus frequencies of 3Hz or faster, there is at least moderate

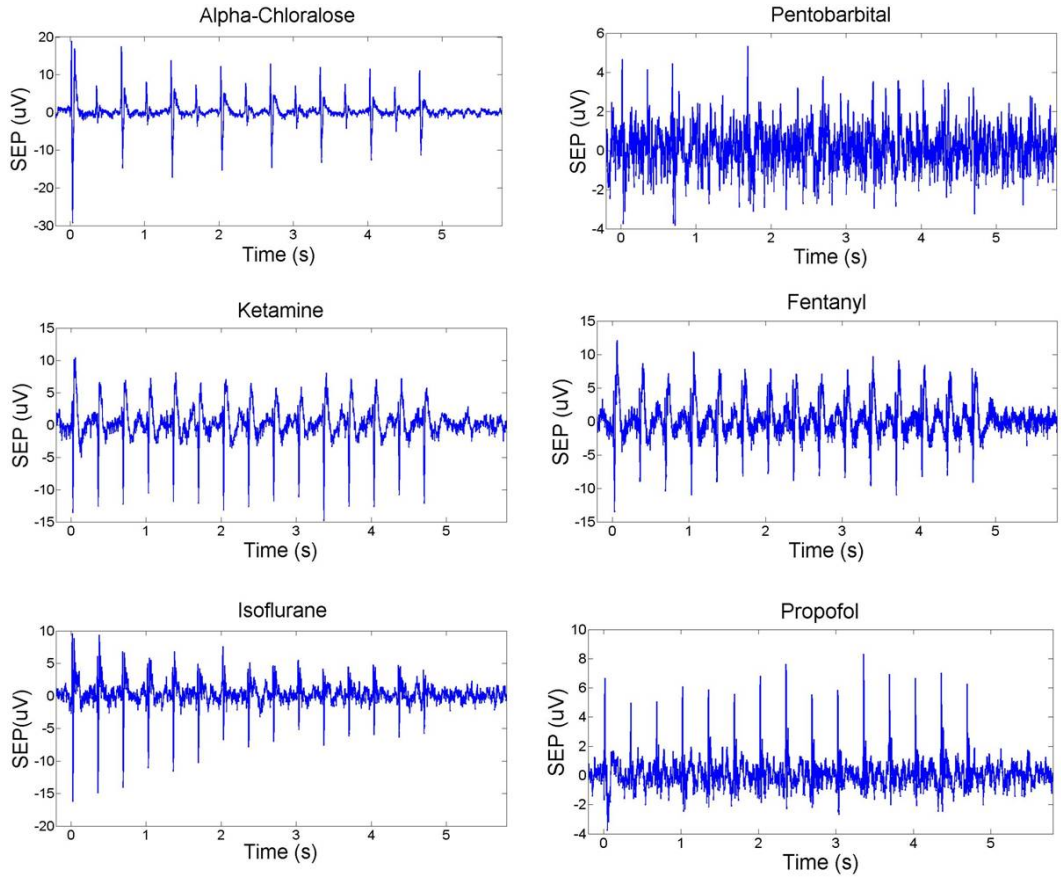

Figure E.3: Average SEP responses under the six anaesthetics to 5s trains of stimuli. Notice that the adaptation is more pronounced for Alpha-Chloralose and Isoflurane than for the other anaesthetics.

adaptation.

- For stimulus frequencies of 6Hz or faster, the adaptation should be strong enough to prevent a subsequent response being discernible from the background signal.

2. Adaptation strength depends on the magnitude of the previous response[1].

These rules can be captured by the following adaptation equation, which calculates the evolution of adaptation to a stimulus occuring at time t_s, i.e. $s(t_s + t)$, using the model fit to the first stimulus in the train, $y(t_1 + t)$ (starting at time t_1 and assumed to last 125ms):

$$\dot{s} = -\frac{1}{\tau_s}(0.5 + s)s + g_s y(t), \tag{E.1}$$

where τ_s and g_s are the time constants and gains of the adaptation process, and the 0.5 imposes the rule that little adaptation should occur when stimuli are delivered at 2Hz or slower.

Use of the same state variable s over the full stimulus train, along with the constraint imposed by Equ.E.1 that s $\geq$ 0, ensures that adaptation is cumulative over stimulation trains. In addition, Equ.E.1 models a second process through its first term, namely a forgetting process that diminishes the cumulative effect while no stimulus is applied.

In their work, Laxminarayan et al. first reduced the number of free parameters to six (corresponding to C, H_e, H_i, R2, τ_e and τ_i in Combined DCM), and then estimated the relationships between these six parameters and s from scatter graphs [92]. As the ultimate aim for their adaptation work was to qualitatively assess the impact of neural adaptation on hemodynamic responses, it was inconsequential to

[1]This phenomenon explains why an alternating pattern of response magnitude can occur (like it does for the Alpha-Chloralose responses in the anaesthetics dataset, see Fig.E.3): as the response to the first stimulation is large (because no adaptation is occuring yet), the response to the second response is highly attenuated, meaning that the third response experiences very little attenuation, etc.

their approach that modelling neural adaptation afforded very little insight about the biological mechanisms of adaptation. For this reason, their approach would have been adequate for simulating hemodynamics from anaesthetic-specific neuronal activity exhibiting adaptation in Chapter 3, but it would not have been sufficient for identifying which specific biological effects of the anaesthetics mask significant neuronal activity in Chapter 4. A more suitable method for integrating neural adaptation dynamics into Combined DCM could be:

1. Fit DCM-EEG to the response to the first stimulus in the train and use this to simulate the (same) response to each stimulation in the train.
2. Use Variational Bayes to optimise the free parameters in $s(t)$ so that the resulting simulated responses fit to the true train of responses. As the adaptation rules captured by Equ.E.1 are specific to one dataset (collected under Alpha-Chloralose anaesthesia), the 0.5 would need to be replaced by a third free parameter as the first rule (there is little or no adaptation when stimuli are delivered at 2Hz or slower) may not be true for other datasets.

E.3 Capturing the post-stimulus undershoot more correctly

It has already been mentioned that Combined DCM is able to capture the initial dip occurring in some of the [ΔHbR] responses in the anaesthetics dataset. A second transient feature, called the post-stimulus undershoot, is also visible in some of the responses, as can be seen in Fig.E.4.

Combined DCM is able to capture this transient because DCM-fMRI models the vasodilatory signal as a damped simple harmonic oscillator:

$$\ddot{s} = x - \kappa\dot{s} - \gamma(s - 1). \tag{E.2}$$

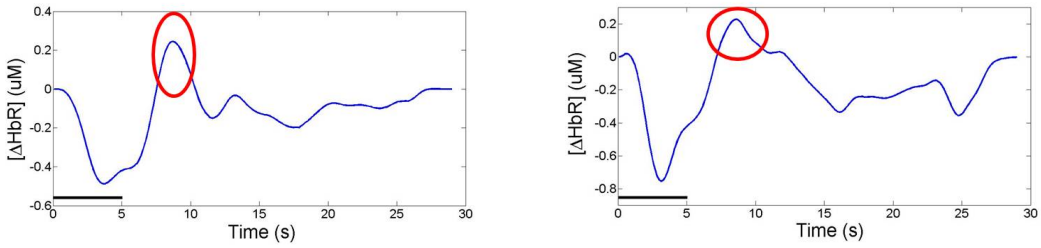

(a) (b)

Figure E.4: Average [ΔHbR] responses under (a) Ketamine and (b) Fentanyl. Post-stimulus undershoots are circled in red.

The decision to model the signal in this way was not experimentally motivated, however [61]. Because hypotheses for the origin of the post-stimulus undershoot have now been developed from experimental results, Combined DCM would be improved by removing the oscillatory dynamics of the vasodilatory signal and incorporating one or more of these hypothesised mechanisms.

One hypothesised mechanism, supported by the findings of Lu et al., 2004 [105], is already captured by Combined DCM, namely uncoupling of oxygen metabolism and CBF. This enables oxygen metabolism to persist when CBF is reducing to resting-state levels.

However, the widespread belief is that the main origin of the post-stimulus undershoot is uncoupling of CBF and CBV, resulting from blood vessels being viscoelastic rather than only elastic [109]. This is not currently captured in Combined DCM.

A number of modified Balloon Models have been developed that extend Grubb's law, which was derived from steady-state experiments, in order to capture the viscoelasticity of blood vessels observed during CBF changes [153]. The effects of this viscoelasticity can be seen in a pressure-volume plot of an active blood vessel: because $\dot{v}$ is different during dilation and constriction, and because the circumferential stress relaxation of blood vessels is dependent on $\dot{v}$, the P-V trajectory follows a different path during dilation and constriction.

The earliest attempts to incorporate viscoelastic effects into the Balloon Model came from Mandeville et al. and Buxton et al. [109][24]. Mandeville et al. added a new variable, A, which was switched on for a short time after stimulation onset and cessation, in order to model vascular tone as an exponential function of time. However Buxton et al., and later Zheng et al., developed time-invariant versions that removed the possibility for discontinuities to occur in the f_{out}-v space by making the relationship a function of $\dot{v}$ [24][203]:

- Buxton et al.: $f_{out} = v^{\frac{1}{\alpha}} + \tau_v v^{-\frac{1}{2}} \dot{v}$[1]
- Zheng et al: $f_{out} = \frac{v^{\frac{1}{\alpha}}}{w}$, where w is a state variable representing vascular tone, modelled by $\tau_w \dot{w} + w = e^{-b\dot{v}}$

The post-stimulus undershoot cannot be captured in [ΔHbR] signals by either of these time-invariant models because it can only occur when $q > v$, i.e. when $m(t) > v^{\frac{1}{\alpha}} + \tau\dot{q}$. When oxygen metabolism is coupled with CBF, i.e. $m(t) = \frac{f(1-(1-E0)^{\frac{1}{f}})}{E0}$, scale analysis shows that this is only possible when $f < \dot{v}$, but neither of the models allow this as it would require $f < \frac{f - v^{\frac{1}{\alpha}}}{\tau + \tau_v v^{-\frac{1}{2}}}$ and $f < \frac{1}{\tau}(f - \frac{v^{\frac{1}{\alpha}}}{w})$, respectively. When oxygen metabolism and CBF are uncoupled, however, there is only the requirement that $m(t) > v^{\frac{1}{\alpha}} + \tau\dot{q}$, which is plausible. Therefore, in order for viscoelastic effects to be able to effect a post-stimulus undershoot in [ΔHbR] signals, it is necessary for both CBF-CBV and oxygen metabolism-CBF to be uncoupled.

E.4 Reducing model complexity

So far, suggestions for how to improve Combined DCM have only involved adding more parameters. Given that the aim should always be to remove model complexity with insufficient modelling power, if Combined DCM were fitted to

[1] Buxton et al. altered this model a few years later to allow τ_v to be different during dilation and constriction [26], however this makes the model time-variant, like Mandeville et al.'s model, as a switching function must be used to detect when v changes direction.

[ΔHbR] data again, α ought to be fixed to its prior value because it has the same effect on v and q, meaning that its influence cancels out in [ΔHbR]. It is important to understand that this modification to Combined DCM is not generalisable to other hemodynamics imaging modalities as α has a large effect on the f_{out}-v space that, for example, the BOLD signal is sensitive to.

Conversely, both r and E0 only control [ΔHbR] response amplitude, meaning that their unique contributions to the [ΔHbR] response cannot be disambiguated. Because these parameters reflect two biologically realistic stages where signal amplitude is modulated during neurovascular coupling, it is necessary to keep both parameters and simply acknowledge that conclusions cannot be made about either stage of amplitude modulation from their individual posterior values.

References

[1] *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. Academic Press, 1 edition, December 2006. 185

[2] M. Abeles. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, 1991. 53

[3] G. K. Aguirre, E. Zarahn, and M. D'esposito. The variability of human, BOLD hemodynamic responses. *NeuroImage*, 8(4):360–369, November 1998. 76

[4] Beau M. Ances, Oleg Leontiev, Joanna E. Perthen, Christine Liang, Amy E. Lansing, and Richard B. Buxton. Regional differences in the coupling of cerebral blood flow and oxygen metabolism changes in response to activation: Implications for BOLD-fMRI. *NeuroImage*, 39(4):1510–1521, February 2008. 2

[5] A. A. Artru, Y. Shapira, and T. A. Bowdle. Electroencephalogram, cerebral metabolic, and vascular responses to propofol anesthesia in dogs. *Journal of neurosurgical anesthesiology*, 4(2):99–109, April 1992. 174

[6] John Ashburner, Gareth Barnes, Chun-Chuan Chen, Jean Daunizeau, Guillaume Flandin, Karl Friston, Stefan Kiebel, James Kilner, Vladimir Litvak, Rosalyn Moran, Will Penny, Maria Rosa, Klaas Stephan, Darren Gitelman, Rik Henson, Chloe Hutton, Volkmar Glauche, Jeremie Mattout,

and Christophe Phillips. *SPM8 Manual*. Wellcome Trust Centre for Neuroimaging Institute of Neurology, UCL, August 2011. 42

[7] AstraZeneca. Study Where Pharmaco Magnetic Resonance Imaging (MRI) Effects of AZD6765 Will be Compared to Placebo in Depressive Male and Female Subjects. Technical report, http://clinicaltrials.gov/ct2/show/NCT01046630, December 2011. 163

[8] David Attwell, Alastair M. Buchan, Serge Charpak, Martin Lauritzen, Brian A. MacVicar, and Eric A. Newman. Glial and neuronal control of brain blood flow. *Nature*, 468(7321):232–243, November 2010. 5, 6, 49

[9] L. Baumane, M. Dzintare, L. Zvejniece, D. Meirena, L. Lauberte, V. Sile, I. Kalvinsh, and N. Sjakste. Increased synthesis of nitric oxide in rat brain cortex due to halogenated volatile anesthetics confirmed by EPR spectroscopy. *Acta Anaesthesiologica Scandinavica*, 46(4):378–383, April 2002. 174

[10] M. M. Bednar and C. E. Gross. Aspirin reduces experimental cerebral blood flow in vivo. *Neurological research*, 21(5):488–490, July 1999. 169

[11] RobertD Bell and BerislavV Zlokovic. Neurovascular mechanisms and bloodbrain barrier disorder in Alzheimer's disease. *Acta neuropathologica*, 118(1):103–113, July 2009. 108

[12] J. Berg-Johnsen and I. A. Langmoen. The effect of isoflurane on excitatory synaptic transmission in the rat hippocampus. *Acta anaesthesiologica Scandinavica*, 36(4):350–355, May 1992. 174, 195

[13] Melanie Boly, Marta Isabel I. Garrido, Olivia Gosseries, Marie-Aurélie A. Bruno, Pierre Boveroux, Caroline Schnakers, Marcello Massimini, Vladimir Litvak, Steven Laureys, and Karl Friston. Preserved feedforward but

impaired top-down processes in the vegetative state. *Science (New York, N.Y.)*, 332(6031):858–862, May 2011. 41

[14] Gilles Bonvento, Robert Charbonné, Jean-Loup Corrèze, Josiane Borredon, Jacques Seylaz, and Pierre Lacombe. Is -chloralose plus halothane induction a suitable anesthetic regimen for cerebrovascular research? *Brain Research*, 665(2):213–221, December 1994. 173

[15] David Borsook, David Bleakman, Richard Hargreaves, Jaymin Upadhyay, Karl F. Schmidt, and Lino Becerra. A 'BOLD' experiment in defining the utility of fMRI in drug development. *NeuroImage*, 42(2):461–466, August 2008. 162

[16] Geoffrey M. Boynton, Stephen A. Engel, Gary H. Glover, and David J. Heeger. Linear Systems Analysis of Functional Magnetic Resonance Imaging in Human V1. *The Journal of Neuroscience*, 16(13):4207–4221, July 1996. 8, 87

[17] Kay H. Brodersen, Lorenz Deserno, Florian Schlagenhauf, Zhihao Lin, Will D. Penny, Joachim M. Buhmann, and Klaas E. Stephan. Dissecting psychiatric spectrum disorders by generative embedding. *NeuroImage: Clinical*, 4:98–111, January 2014. 143

[18] Kay H. Brodersen, Florent Haiss, Cheng S. Ong, Fabienne Jung, Marc Tittgemeyer, Joachim M. Buhmann, Bruno Weber, and Klaas E. Stephan. Model-based feature construction for multivariate decoding. *NeuroImage*, 56(2):601–615, May 2011. 9, 41, 105

[19] Kay H. Brodersen, Thomas M. Schofield, Alexander P. Leff, Cheng Soon S. Ong, Ekaterina I. Lomakina, Joachim M. Buhmann, and Klaas E. Stephan. Generative embedding for model-based classification of fMRI data. *PLoS computational biology*, 7(6):e1002079+, June 2011. 9, 41, 105

[20] Ed Bullmore. The future of functional MRI in clinical medicine. *NeuroImage*, 62(2):1267–1271, August 2012. 158, 161

[21] Edward Bullmore, Michael Brammer, Steve C. R. Williams, Sophia Rabe-Hesketh, Nicolas Janot, Anthony David, John Mellers, Robert Howard, and Pak Sham. Statistical methods of estimation and inference for functional MR image analysis. *Magn. Reson. Med.*, 35(2):261–277, February 1996. 182

[22] Broc A. Burke and Solomon G. Diamond. Measuring cerebral hemodynamics with a modified magnetoencephalography system. *Physiological measurement*, 33(12):2079–2098, December 2012. 178

[23] David W. Busija, Ferenc Bari, Ferenc Domoki, and Thomas Louis. Mechanisms involved in the cerebrovascular dilator effects of N-methyl-d-aspartate in cerebral cortex. *Brain research reviews*, 56(1):89–100, November 2007. 11

[24] R. B. Buxton, K. Miller, L. R. Frank, and E. C. Wong. BOLD signal dynamics: the balloon model with viscoelastic effects. *Proc. Int. Mag. Reson. Med*, 6:1401+, 1998. 202

[25] R. B. Buxton, E. C. Wong, and L. R. Frank. Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. *Magnetic resonance in medicine*, 39(6):855–864, June 1998. 46

[26] Richard B. Buxton, Kâmil Uludağ, David J. Dubowitz, and Thomas T. Liu. Modeling the hemodynamic response to brain activation. *NeuroImage*, 23:S220–S233, January 2004. 189, 202

[27] Gyorgy Buzsaki, Costas A. Anastassiou, and Christof Koch. The origin of extracellular fields and currents â EEG, ECoG, LFP and spikes. *Nat Rev Neurosci*, 13(6):407–420, May 2012. 46

[28] William H. Calvin and George A. Ojemann. Conversations with Neil's brain. 1994. 24

[29] C. Carlsson, D. S. Smith, M. M. Keykhah, I. Englebach, and J. R. Harp. The effects of high-dose fentanyl on cerebral circulation and metabolism in rats. *Anesthesiology*, 57(5):375–380, November 1982. 173

[30] Milena Cavazzuti, Carlo A. Porro, Gian Paolo Biral, Carlo Benassi, and Gian Carlo Barbieri. Ketamine Effects on Local Cerebral Blood Flow and Metabolism in the Rat. *Journal of Cerebral Blood Flow & Metabolism*, 7(6):806–811, December 1987. 173

[31] Experimental MRI Centre. MRI Basics. 4

[32] Rajnish K. Chaturvedi and M. Flint Beal. Mitochondrial Diseases of the Brain. *Free Radical Biology and Medicine*, April 2013. 108

[33] J. J. Claus, M. M. B. Breteler, D. Hasan, E. P. Krenning, M. L. Bots, D. E. Grobbee, J. C. Van Swieten, F. Van Harskamp, and A. Hofman. Regional Cerebral Blood Flow and Cerebrovascular Risk Factors in the Elderly Population. *Neurobiology of Aging*, 19(1):57–64, January 1998. 169

[34] CarolA Colton. Heterogeneity of Microglial Activation in the Innate Immune Response in the Brain. *Journal of Neuroimmune Pharmacology*, 4(4):399–418, 2009. 113

[35] Ghana B. Corporation. Four hospitals receive MRI machines. January 2012. 160

[36] Yvonne Couch, Chris J. Martin, Clare Howarth, Josie Raley, Alexandre A. Khrapitchev, Michael Stratford, Trevor Sharp, Nicola R. Sibson, and Daniel C. Anthony. Systemic inflammation alters central 5-HT function as determined by pharmacological MRI. *NeuroImage*, 75:177–186, July 2013. 176

[37] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. John Wiley and Sons Inc., Hoboken, New Jersey, USA, 2006. 125

[38] G. Crosby, A. M. Crane, and L. Sokoloff. Local changes in cerebral glucose utilization during ketamine anesthesia. *Anesthesiology*, 56(6):437–443, June 1982. 173

[39] D. R. Curtis and D. Lodge. Pentobarbitone enhancement of the inhibitory action of GABA. *Nature*, 270(5637):543–544, December 1977. 173

[40] J. Daut, W. Maier-Rudolph, N. von Beckerath, G. Mehrke, K. Gunther, and L. Goedel-Meinen. Hypoxic dilation of coronary arteries is mediated by ATP-sensitive potassium channels. *Science*, 247(4948):1341–1344, March 1990. 2

[41] Olivier David, James M. Kilner, and Karl J. Friston. Mechanisms of evoked and induced responses in MEG/EEG. *NeuroImage*, 31(4):1580–1591, July 2006. 41, 44, 184

[42] Mark D'Esposito, Leon Y. Deouell, and Adam Gazzaley. Alterations in the BOLD fMRI signal with ageing and disease: a challenge for neuroimaging. *Nature Reviews Neuroscience*, 4(11):863–872, November 2003. 2, 7, 169

[43] Mark D'Esposito, Eric Zarahn, Geoffrey K. Aguirre, and Bart Rypma. The Effect of Normal Aging on the Coupling of Neural Activity to the Bold Hemodynamic Response. *NeuroImage*, 10(1):6–14, July 1999. 169

[44] Anna Devor, Elizabeth M. Hillman, Peifang Tian, Christian Waeber, Ivan C. Teng, Lana Ruvinskaya, Mark H. Shalinsky, Haihao Zhu, Robert H. Haslinger, Suresh N. Narayanan, Istvan Ulbert, Andrew K. Dunn, Eng H. Lo, Bruce R. Rosen, Anders M. Dale, David Kleinfeld, and David A. Boas. Stimulus-induced changes in blood flow and 2-deoxyglucose uptake dissociate in ipsilateral somatosensory cortex. *The Journal of neuroscience : the*

official journal of the Society for Neuroscience, 28(53):14347–14357, December 2008. 5

[45] Anna Devor, Sava Sakadzic, Vivek J. Srinivasan, Mohammad A. Yaseen, Krystal Nizar, Payam A. Saisan, Peifang Tian, Anders M. Dale, Sergei A. Vinogradov, Maria A. Franceschini, and David A. Boas. Frontiers in optical imaging of cerebral blood flow and metabolism. *Journal of Cerebral Blood Flow & Metabolism*, 32(7):1259–1276, January 2012. 9

[46] J. C. Drummond and M. M. Todd. The response of the feline cerebral circulation to PaCO2 during anesthesia with isoflurane and halothane and during sedation with nitrous oxide. *Anesthesiology*, 62(3):268–273, March 1985. 174

[47] Robert E. Dudley, Stanley R. Nelson, and Fred Samson. Influence of chloralose on brain regional glucose utilization. *Brain Research*, 233(1):173–180, February 1982. 173

[48] Jeff H. Duyn. The future of ultra-high field MRI and fMRI for study of the human brain. *NeuroImage*, 62(2):1241–1248, August 2012. 160

[49] Clark N. Eid and Robert J. Gorlin. Cobalt-Induced Polycythemia. *Journal of Dental Research*, 37(6):1141–1148, November 1958. 47

[50] C. Eng, A. M. Lam, T. S. Mayberg, C. Lee, and T. Mathisen. The influence of propofol with and without nitrous oxide on cerebral blood flow velocity and CO2 reactivity in humans. *Anesthesiology*, 77(5):872–879, November 1992. 174

[51] Jutta Engel, Howard A. Schultens, and Detlev Schild. Small Conductance Potassium Channels Cause an Activity-Dependent Spike Frequency Adaptation and Make the Transfer Function of Neurons Logarithmic. *Biophysical Journal*, 76(3):1310–1319, March 1999. 196

[52] Siamac Fazli, Jan Mehnert, Jens Steinbrink, Gabriel Curio, Arno Villringer, Klaus-Robert Müller, and Benjamin Blankertz. Enhanced performance by a hybrid NIRSEEG brain computer interface. *NeuroImage*, 59(1):519–529, January 2012. 151

[53] Anna Fiorentino. A Thinking Cap, April 2012. 104

[54] Organisation for Economic Co-operation and Development. OECD Health Data 2012. Technical report, http://www.oecd.org/health/health-systems/oecdhealthdata2012.htm, June 2012. 160

[55] Maria A. Franceschini, Ilkka Nissilä, Weicheng Wu, Solomon G. Diamond, Giorgio Bonmassar, and David A. Boas. Coupling between somatosensory evoked potentials and hemodynamic response in the rat. *NeuroImage*, 41(2):189–203, June 2008. 197

[56] Maria Angela A. Franceschini, Harsha Radhakrishnan, Kiran Thakur, Weicheng Wu, Svetlana Ruvinskaya, Stefan Carp, and David A. Boas. The effect of different anesthetics on neurovascular coupling. *NeuroImage*, 51(4):1367–1377, July 2010. 6, 10, 19, 179

[57] Giovanni B. Frisoni, Nick C. Fox, Clifford R. Jack, Philip Scheltens, and Paul M. Thompson. The clinical use of structural MRI in Alzheimer disease. *Nature reviews. Neurology*, 6(2):67–77, February 2010. 106

[58] K. J. Friston. Bayesian estimation of dynamical systems: an application to fMRI. *Neuroimage*, 16(2):513–530, June 2002. 56, 187

[59] K. J. Friston, L. Harrison, and W. Penny. Dynamic causal modelling. *NeuroImage*, 19(4):1273–1302, August 2003. 41, 47, 56, 187

[60] K. J. Friston, P. Jezzard, and R. Turner. Analysis of functional MRI time-series. *Hum. Brain Mapp.*, 1(2):153–171, 1994. 24

[61] K. J. Friston, A. Mechelli, R. Turner, and C. J. Price. Nonlinear Responses in fMRI: The Balloon Model, Volterra Kernels, and Other Hemodynamics. *NeuroImage*, 12(4):466–477, October 2000. 87, 201

[62] Karl J. Friston, Oliver Josephs, Geraint Rees, and Robert Turner. Nonlinear event-related responses in fMRI. *Magn. Reson. Med.*, 39(1):41–52, January 1998. 27, 74, 89

[63] V. R. Fuchs and H. C. Sox. Physicians' views of the relative importance of thirty medical innovations. *Health affairs (Project Hope)*, 20(5):30–42, 2001. 154

[64] I. Fulga and T. W. Stone. Comparison of an adenosine A1 receptor agonist and antagonist on the rat EEG. *Neuroscience Letters*, 244(1):55–59, March 1998. 11

[65] Kennon M. Garrett and Jiangping Gan. Enhancement of -Aminobutyric AcidA Receptor Activity by -Chloralose. *Journal of Pharmacology and Experimental Therapeutics*, 285(2):680–686, May 1998. 173

[66] T. J. Gawne and B. J. Richmond. How independent are the messages carried by adjacent inferior temporal cortical neurons? *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 13(7):2758–2771, July 1993. 125

[67] Paolo Giacometti and Solomon G. Diamond. Compliant head probe for positioning electroencephalography electrodes and near-infrared spectroscopy optodes. *Journal of Biomedical Optics*, 18(2):027005, 2013. 14, 104

[68] Timo Grimmer, Matthias Riemenschneider, Hans Förstl, Gjermund Henriksen, William E. Klunk, Chester A. Mathis, Tohru Shiga, Hans-Jürgen Wester, Alexander Kurz, and Alexander Drzezga. Beta Amyloid in Alzheimer's Disease: Increased Deposition in Brain Is Reflected in Reduced

Concentration in Cerebrospinal Fluid. *Biological Psychiatry*, 65(11):927–934, June 2009. 106

[69] Tim Hahn, Sebastian Heinzel, Michael M. Plichta, Andreas Reif, Klaus-Peter Lesch, and Andreas J. Fallgatter. Neurovascular Coupling in the Human Visual Cortex Is Modulated by Cyclooxygenase-1 (COX-1) Gene Variant. *Cerebral Cortex*, 21(7):1659–1666, July 2011. 109

[70] David Hailey. CT and MRI Services in Alberta: Comparisons With Other Health Care Systems. Technical report, http://www.ihe.ca/publications/library/2008/ct-and-mri-services-in-alberta-comparisons-with-other-health-care-systems/, 2008. 163

[71] Tim G. Hales and Jeremy J. Lambert. The actions of propofol on inhibitory amino acid receptors of bovine adrenomedullary chromaffin cells and rodent central neurones. *British Journal of Pharmacology*, 104(3):619–628, November 1991. 174

[72] Daniel A. Handwerker, John M. Ollinger, and Mark D'Esposito. Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *NeuroImage*, 21(4):1639–1651, April 2004. 2, 17

[73] John-Dylan Haynes and Geraint Rees. Decoding mental states from brain activity in humans. *Nature Reviews Neuroscience*, 7(7):523–534, July 2006. 104

[74] D. J. Heeger, A. C. Huk, W. S. Geisler, and D. G. Albrecht. Spikes versus BOLD: what does neuroimaging tell us about neuronal activity? *Nature neuroscience*, 3(7):631–633, July 2000. 51

[75] Kristy S. Hendrich, Patrick M. Kochanek, John A. Melick, Joanne K. Schiding, Kimberly D. Statler, Donald S. Williams, Donald W. Marion, and Chien Ho. Cerebral perfusion during anesthesia with fentanyl, isoflurane, or pentobarbital in normal rats studied by arterial spin-labeled MRI. *Magn. Reson. Med.*, 46(1):202–206, July 2001. 173, 174

[76] Harald Hentschke, Cornelius Schwarz, and Bernd Antkowiak. Neocortex is the major target of sedative concentrations of volatile anaesthetics: strong depression of firing rates and increase of GABAA receptor-mediated inhibition. *European Journal of Neuroscience*, 21(1):93–102, January 2005. 174

[77] A. L. HODGKIN and A. F. HUXLEY. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of physiology*, 117(4):500–544, August 1952. 195

[78] Great Ormond Street Hospital. Functional magnetic resonance imaging (fMRI). 2013. 161

[79] Gholam-Ali Hossein-Zadeh, Babak A. Ardekani, and Hamid Soltanian-Zadeh. A signal subspace approach for modeling the hemodynamic response function in fMRI. *Magnetic Resonance Imaging*, 21(8):835–843, October 2003. 76

[80] S. A. Huettel. Non-linearities in the blood-oxygenation-level dependent (BOLD) response measured by functional magnetic resonance imaging (fMRI). In *Engineering in Medicine and Biology Society, 2004. IEMBS '04. 26th Annual International Conference of the IEEE*, volume 2, pages 4413–4416. IEEE, 2004. 80

[81] Scott A. Huettel and Gregory McCarthy. Evidence for a Refractory Period in the Hemodynamic Response to Visual Stimuli as Measured by MRI. *NeuroImage*, 11(5):547–553, May 2000. 80

[82] Costantino Iadecola. Neurovascular regulation in the normal brain and in Alzheimer's disease. *Nature Reviews Neuroscience*, 5(5):347–360, May 2004. 107, 108, 109

[83] Clifford R. Jack, David S. Knopman, William J. Jagust, Ronald C. Petersen, Michael W. Weiner, Paul S. Aisen, Leslie M. Shaw, Prashanthi Vemuri, Heather J. Wiste, Stephen D. Weigand, Timothy G. Lesnick, Vernon S. Pankratz, Michael C. Donohue, and John Q. Trojanowski. Tracking pathophysiological processes in Alzheimer's disease: an updated hypothetical model of dynamic biomarkers. *Lancet neurology*, 12(2):207–216, February 2013. 107

[84] B. H. Jansen and V. G. Rit. Electroencephalogram and visual evoked potential generation in a mathematical model of coupled cortical columns. *Biological cybernetics*, 73(4):357–366, September 1995. 42, 56

[85] Karl A. Kasischke, Harshad D. Vishwasrao, Patricia J. Fisher, Warren R. Zipfel, and Watt W. Webb. Neural activity triggers neuronal oxidative metabolism followed by astrocytic glycolysis. *Science (New York, N.Y.)*, 305(5680):99–103, July 2004. 5

[86] J. M. Kilner, J. Mattout, R. Henson, and K. J. Friston. Hemodynamic correlates of EEG: A heuristic. *NeuroImage*, 28(1):280–286, October 2005. 50

[87] William E. Klunk, Henry Engler, Agneta Nordberg, Yanming Wang, Gunnar Blomqvist, Daniel P. Holt, Mats Bergström, Irina Savitcheva, Guo-feng F. Huang, Sergio Estrada, Birgitta Ausén, Manik L. Debnath, Julien Barletta, Julie C. Price, Johan Sandell, Brian J. Lopresti, Anders Wall, Pernilla Koivisto, Gunnar Antoni, Chester A. Mathis, and Bengt Långström. Imaging brain amyloid in Alzheimer's disease with Pittsburgh Compound-B. *Annals of neurology*, 55(3):306–319, March 2004. 106

[88] K. R. Ko, A. C. Ngai, and H. R. Winn. Role of adenosine in regulation of regional cerebral blood flow in sensory cortex. *The American journal of physiology*, 259(6 Pt 2), December 1990. 11

[89] E. Kouvaras, E. K. Asprodini, I. Asouchidou, A. Vasilaki, T. Kilindris, D. Michaloudis, I. Koukoutianou, C. Papatheodoropoulos, and G. Kostopoulos. Fentanyl treatment reduces GABAergic inhibition in the CA1 area of the hippocampus 24h after acute exposure to the drug. *Neuropharmacology*, 55(7):1172–1182, December 2008. 173

[90] L. F. Kozachenko and N. N. Leonenko. Sample Estimate of the Entropy of a Random Vector. *Probl. Peredachi Inf.*, 23(2):9–16, 1987. 129

[91] Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual information. *Physical Review E*, 69(6):066138+, June 2004. 129, 131, 137

[92] Srinivas Laxminarayan, Gilead Tadmor, Solomon G. Diamond, Eric Miller, Maria Angela, and Dana H. Brooks. Modeling habituation in rat EEG-evoked responses via a neural mass model with feedback. *Biological cybernetics*, 105(5-6):371–397, December 2011. 199

[93] Andrea Lazosky, G. Bryan Young, Sandra Zirul, and Renee Phillips. Quality of life after septic illness. *Journal of critical care*, 25(3):406–412, September 2010. 122

[94] Joseph G. Lee, Antal G. Hudetz, Jeremy J. Smith, Cecilia J. Hillard, Zeljko J. Bosnjak, and John P. Kampine. The Effects of Halothane and Isoflurane on Cerebrocortical Microcirculation and Autoregulation as Assessed by LaserDoppler Flowmetry. *Anesthesia & Analgesia*, 79(1):58–65, July 1994. 174

[95] Hao Lei, O. Grinberg, C. I. Nwaigwe, H. G. Hou, H. Williams, H. M. Swartz, and J. F. Dunn. The effects of ketaminexylazine anesthesia on cerebral blood flow and oxygenation observed using nuclear magnetic resonance perfusion imaging and electron paramagnetic resonance oximetry. *Brain Research*, 913(2):174–179, September 2001. 173

[96] Steven C. Leiser, John Dunlop, Mark R. Bowlby, and David M. Devilbiss. Aligning strategies for using EEG as a surrogate biomarker: a review of preclinical and clinical research. *Biochemical pharmacology*, 81(12):1408–1421, June 2011. 162

[97] J. Li and C. Iadecola. Nitric oxide and adenosine mediate vasodilation during functional activation in cerebellar cortex. *Neuropharmacology*, 33(11):1453–1461, November 1994. 51

[98] Chih-Kai K. Liao, Seu-Mei M. Wang, Yuh-Lien L. Chen, Hwai-Shi S. Wang, and Jiahn-Chun C. Wu. Lipopolysaccharide-induced inhibition of connexin43 gap junction communication in astrocytes is mediated by downregulation of caveolin-3. *The international journal of biochemistry & cell biology*, 42(5):762–770, May 2010. 109

[99] L. H. Lin, L. L. Chen, J. A. Zirrolli, and R. A. Harris. General anesthetics potentiate gamma-aminobutyric acid actions on gamma-aminobutyric acidA receptors expressed by Xenopus oocytes: lack of involvement of intracellular calcium. *Journal of Pharmacology and Experimental Therapeutics*, 263(2):569–578, November 1992. 173

[100] Martin A. Lindquist, Ji Meng Loh, Lauren Y. Atlas, and Tor D. Wager. Modeling the hemodynamic response function in fMRI: Efficiency, bias and mis-modeling. *NeuroImage*, 45(1):S187–S198, March 2009. 26

[101] Martin A. Lindquist and Tor D. Wager. Validity and power in hemodynamic

response modeling: A comparison study and a new approach. *Hum. Brain Mapp.*, 28(8):764–784, August 2007. 17

[102] D. Lodge and N. A. Anis. Effects of ketamine and three other anaesthetics on spinal reflexes and inhibitions in the cat. *British Journal of Anaesthesia*, 56(10):1143–1152, October 1984. 173

[103] Nikos K. Logothetis. What we can do and what we cannot do with fMRI. *Nature*, 453(7197):869–878, June 2008. 3

[104] Nikos K. Logothetis, Jon Pauls, Mark Augath, Torsten Trinath, and Axel Oeltermann. Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412(6843):150–157, July 2001. 8, 50, 87

[105] Hanzhang Lu, Xavier Golay, James J. Pekar, and Peter C. M. van Zijl. Sustained Poststimulus Elevation in Cerebral Oxygen Utilization After Vascular Recovery. *Journal of Cerebral Blood Flow & Metabolism*, 24(7):764–770, July 2004. 201

[106] T. Maekawa, C. Tommasino, H. M. Shapiro, J. Keifer-Goodman, and R. W. Kohlenberger. Local cerebral blood flow and glucose utilization during isoflurane anesthesia in the rat. *Anesthesiology*, 65(2):144–151, August 1986. 174

[107] PierreJ Magistretti and Igor Allaman. Brain Energy Metabolism. In DonaldW Pfaff, editor, *Neuroscience in the 21st Century*, pages 1591–1620. Springer New York, 2013. 5

[108] Dov Malonek, Ulrich Dirnagl, Ute Lindauer, Katsuya Yamada, Iwao Kanno, and Amiram Grinvald. Vascular imprints of neuronal activity: Relationships between the dynamics of cortical blood flow, oxygenation, and volume changes following sensorystimulation. *Proceedings of the National Academy of Sciences*, 94(26):14826–14831, December 1997. 189

[109] Joseph B. Mandeville, John J. A. Marota, C. Ayata, Greg Zaharchuk, Michael A. Moskowitz, Bruce R. Rosen, and Robert M. Weisskoff. Evidence of a Cerebrovascular Postarteriole Windkessel With Delayed Compliance. *Journal of Cerebral Blood Flow & Metabolism*, 19(6):679–689, June 1999. 201, 202

[110] Chris Martin, John Martindale, Jason Berwick, and John Mayhew. Investigating neuralhemodynamic coupling and the hemodynamic response function in the awake rat. *NeuroImage*, 32(1):33–48, August 2006. 19, 27

[111] Chris Martin and Nicola R. Sibson. Pharmacological MRI in animal models: A useful tool for 5-HT research? *Neuropharmacology*, 55(6):1038–1047, November 2008. 169

[112] Kazuto Masamoto and Iwao Kanno. Anesthesia and the quantitative evaluation of neurovascular coupling. *J Cereb Blood Flow Metab*, 32(7):1233–1247, July 2012. 176

[113] Kazuto Masamoto, Tae Kim, Mitsuhiro Fukuda, Ping Wang, and Seong-Gi G. Kim. Relationship between neural, vascular, and BOLD signals in isoflurane-anesthetized rat somatosensory cortex. *Cerebral cortex (New York, N.Y. : 1991)*, 17(4):942–950, April 2007. 197

[114] Claus Mathiesen, Kirsten Caesar, Nuran Akgören, and Martin Lauritzen. Modification of activity-dependent increases of cerebral blood flow by excitatory synaptic activity and spikes in rat cerebellar cortex. *The Journal of Physiology*, 512(2):555–566, October 1998. 8, 51

[115] Paul M. Matthews. MRI Studies in the Pharmaceutical Industry: Turning Data into Information. 2008. 163

[116] Andrea Mechelli, Cathy J. Price, and Karl J. Friston. Nonlinear Coupling

between Evoked rCBF and BOLD Signals: A Simulation Study of Hemodynamic Responses. *NeuroImage*, 14(4):862–872, October 2001. 80

[117] G. A. Miller. *Note on the bias on information estimates*, pages 95–100. 1955. 192

[118] Rosalyn J. Moran, Nicolas Mallet, Vladimir Litvak, Raymond J. Dolan, Peter J. Magill, Karl J. Friston, and Peter Brown. Alterations in brain connectivity underlying beta oscillations in parkinsonism. *PLoS computational biology*, 7(8), August 2011. 11, 41

[119] Victoria L. Morgan, Benoit M. Dawant, Yong Li, and David R. Pickens. Comparison of fMRI statistical software packages and strategies for analysis of images containing random and stimulus-correlated motion. *Computerized Medical Imaging and Graphics*, 31(6):436–446, September 2007. 27, 182

[120] Janaina Mourão Miranda, Arun L. Bokde, Christine Born, Harald Hampel, and Martin Stetter. Classifying brain states and determining the discriminating activation patterns: Support Vector Machine on functional MRI data. *NeuroImage*, 28(4):980–995, December 2005. 104

[121] Todd A. Mulderink, Darren R. Gitelman, Mesulam, and Todd B. Parrish. On the Use of Caffeine as a Contrast Booster for BOLD fMRI Studies. *NeuroImage*, 15(1):37–44, January 2002. 7, 169

[122] Shingo Murakami and Yoshio Okada. Contributions of principal neocortical neurons to magnetoencephalography and electroencephalography signals. *The Journal of Physiology*, 575(3):925–936, September 2006. 46

[123] Kiyoshi Nagase, Hiroki Iida, and Shuji Dohi. l-Arginine and Nitroglycerin Restore Hypercapnia-Induced Cerebral Vasodilation in Rabbits After its Attenuation by Ketamine. *Anesthesia & Analgesia*, 94(4):954–958, April 2002. 173

[124] Yasuaki Nakao, Yoshiaki Itoh, Tang-Yong Kuang, Michelle Cook, Jane Jehle, and Louis Sokoloff. Effects of anesthesia on functional activation of cerebral blood flow and metabolism. *Proceedings of the National Academy of Sciences*, 98(13):7593–7598, June 2001. 173

[125] Ilya Nemenman, William Bialek, and Rob de Ruyter van Steveninck. Entropy and information in neural spike trains: Progress on the sampling problem. *Physical Review E*, 69(5):056111+, May 2004. 193

[126] A. C. Ngai, M. A. Jolley, R. D'Ambrosio, J. R. Meno, and H. R. Winn. Frequency-dependent changes in cerebral blood flow and evoked potentials during somatosensory stimulation in the rat. *Brain research*, 837(1-2):221–228, August 1999. 6

[127] MitchellS Nobler, KristianR Olvet, and HaroldA Sackeim. Effects of medications on cerebral blood flow in late-life depression. 4(1):51–58, 2002. 169

[128] Takayuki Obata, Thomas T. Liu, Karla L. Miller, Wen-Ming Luh, Eric C. Wong, Lawrence R. Frank, and Richard B. Buxton. Discrepancies between BOLD and flow dynamics in primary and supplementary motor areas: application of the balloon model to the interpretation of BOLD transients. *NeuroImage*, 21(1):144–153, January 2004. 47, 189

[129] C. Obyn and I. Cleemput. The capital cost and productivity of MRI in a Belgian setting. *JBR-BTR : organe de la Société royale belge de radiologie (SRBR) = orgaan van de Koninklijke Belgische Vereniging voor Radiologie (KBVR)*, 93(2):92–96, 2010. 160

[130] Department of Economic, Social Affairs, and Population Division. World Population Ageing: 1950-2050. Technical report,

http://www.un.org/esa/population/publications/worldageing19502050/, 2002. 158

[131] Thayer School of Engineering at Dartmouth. Noninvasive Head Probe for Alzheimer's Disease. Youtube, April 2012. 105

[132] U.S. Department of Health, Human Services, Food, and Drug Administration. Innovation or Stagnation: Challenge and Opportunity on the Critical Path to New Medical Products. Technical report, http://www.fda.gov/ScienceResearch/SpecialTopics/CriticalPathInitiative, March 2004. 154

[133] S. Ogawa, T. M. Lee, A. R. Kay, and D. W. Tank. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences of the United States of America*, 87(24):9868–9872, December 1990. 4

[134] Mike W. Oram, Peter Földiák, David I. Perrett, Mike W. Oram, and Frank Sengpiel. The 'Ideal Homunculus': decoding neural population signals. *Trends in Neurosciences*, 21(6):259–265, June 1998. 125

[135] R. E. Oren, N. A. Rasool, and E. H. Rubinstein. Effect of ketamine on cerebral cortical blood flow and metabolism in rabbits. *Stroke*, 18(2):441–444, March 1987. 173

[136] C. Ori, M. Dam, G. Pizzolato, L. Battistin, and G. Giron. Effects of isoflurane anesthesia on local cerebral glucose utilization in the rat. *Anesthesiology*, 65(2):152–156, August 1986. 174

[137] Peter G. Osborne. Hippocampal and Striatal Blood Flow During Behavior in Rats: Chronic Laser Doppler Flowmetry Study. *Physiology & Behavior*, 61(4):485–492, April 1997. 173

[138] T. Oshima, F. Karasawa, and T. Satoh. Effects of propofol on cerebral blood flow and the metabolic rate of oxygen in humans. *Acta Anaesthesiologica Scandinavica*, 46(7):831–835, August 2002. 174

[139] Dirk Ostwald, Camillo Porcaro, and Andrew P. Bagshaw. An information theoretic approach to EEGfMRI integration of visually evoked responses. *NeuroImage*, 49(1):498–516, January 2010. 125, 194

[140] T. Otsuka, L. Wei, V. R. Acuff, A. Shimizu, K. D. Pettigrew, C. S. Patlak, and J. D. Fenstermacher. Variation in local cerebral blood flow response to high-dose pentobarbital sodium in the rat. *American Journal of Physiology - Heart and Circulatory Physiology*, 261(1):H110–H120, July 1991. 173

[141] T. Otsuka, L. Wei, D. Bereczki, V. Acuff, C. Patlak, and J. Fenstermacher. Pentobarbital produces dissimilar changes in glucose influx and utilization in brain. *American Journal of Physiology - Regulatory, Integrative and Comparative Physiology*, 261(2):R265–R275, August 1991. 173

[142] Liam Paninski. Estimation of Entropy and Mutual Information. *Neural Computation*, 15(6):1191–1253, June 2003. 193

[143] S. Panzeri, S. R. Schultz, A. Treves, and E. T. Rolls. Correlations and the encoding of information in the nervous system. *Proceedings. Biological sciences / The Royal Society*, 266(1423):1001–1012, May 1999. 125

[144] Stefano Panzeri, Cesare Magri, and Nikos K. Logothetis. On the use of information theory for the analysis of the relationship between neural and imaging signals. *Magnetic resonance imaging*, 26(7):1015–1025, September 2008. 125

[145] Stefano Panzeri, Riccardo Senatore, Marcelo A. Montemurro, and Rasmus S. Petersen. Correcting for the Sampling Bias Problem in Spike Train

Information Measures. *Journal of Neurophysiology*, 98(3):1064–1072, September 2007. 127, 192, 193

[146] Stefano Panzeri and Alessandro Treves. Analytical estimates of limited sampling biases in different information measures. *Network: Computation in Neural Systems*, 7:87–107, 1996. 192

[147] Angeliki Papana and Dimitris Kugiumtzis. Evaluation of Mutual Information Estimators for Time Series. April 2009. 126, 143

[148] I. Parker, C. B. Gundersen, and R. Miledi. Actions of pentobarbital on rat brain receptors expressed in Xenopus oocytes. *The Journal of Neuroscience*, 6(8):2290–2297, August 1986. 173

[149] Claire M. Peppiatt, Clare Howarth, Peter Mobbs, and David Attwell. Bidirectional control of CNS capillary diameter by pericytes. *Nature*, 443(7112):700–704, October 2006. 5

[150] Perry. The influence of systemic inflammation on inflammation in the brain: implications for chronic neurodegenerative disease. *Brain, Behavior, and Immunity*, 18(5):407–413, September 2004. 107, 113

[151] Gabor C. Petzold and Venkatesh N. Murthy. Role of Astrocytes in Neurovascular Coupling. *Neuron*, 71(5):782–797, September 2011. 5

[152] Philips. Philips to introduce the first MRI scanner in the public sector in Botswana. May 2011. 160

[153] C. I. Porciuncula, Armstrong, A. C. Guyton, and H. L. Stone. Delayed compliance in external jugular vein of the dog. *American Journal of Physiology – Legacy Content*, 207(3):728–732, September 1964. 201

[154] Scott Prahl. Optical Absorption of Hemoglobin. December 1999. 47

[155] J. Prichard, D. Rothman, E. Novotny, O. Petroff, T. Kuwabara, M. Avison, A. Howseman, C. Hanstock, and R. Shulman. Lactate rise detected by 1H NMR in human visual cortex during physiologic stimulation. *Proceedings of the National Academy of Sciences*, 88(13):5829–5831, July 1991. 53

[156] Martin Prince, Renata Bryce, and Cleusa Ferr. World Alzheimer Report 2011 Alzheimer's Disease International, September 2011. 106

[157] Harsha Radhakrishnan, Weicheng Wu, David Boas, and Maria A. Franceschini. Study of neurovascular coupling by modulating neuronal activity with GABA. *Brain Research*, 1372:1–12, February 2011. 24

[158] Marcus E. Raichle and Debra A. Gusnard. Appraising the brain's energy budget. *Proceedings of the National Academy of Sciences*, 99(16):10237–10239, August 2002. 5

[159] David M. Rector, Robert F. Rogers, James S. Schwaber, Ronald M. Harper, and John S. George. Scattered-Light Imaging in Vivo Tracks Fast and Slow Processes of Neurophysiological Activation. *NeuroImage*, 14(5):977–994, November 2001. 178

[160] L. Q. Ren, N. Gourmala, H. W. Boddeke, and P. J. Gebicke-Haerter. Lipopolysaccharide-induced expression of IP-10 mRNA in rat brain and in cultured rat astrocytes and microglia. *Brain research. Molecular brain research*, 59(2):256–263, August 1998. 114

[161] Jorge J. Riera, Arne Schousboe, Helle S. Waagepetersen, Clare Howarth, and Fahmeed Hyder. The micro-architecture of the cerebral cortex: functional neuroimaging models and metabolism. *NeuroImage*, 40(4):1436–1459, May 2008. 5, 197

[162] Jorge J. Riera, Xiaohong Wan, Juan C. Jimenez, and Ryuta Kawashima.

Nonlinear local electrovascular coupling. I: A theoretical model. *Hum. Brain Mapp.*, 27(11):896–914, November 2006. 50, 51

[163] Maria J. Rosa, James M. Kilner, and Will D. Penny. Bayesian comparison of neurovascular coupling models using EEG-fMRI. *PLoS computational biology*, 7(6), June 2011. 16

[164] Bernhard Rosengarten, Mirko Hecht, Dieter Auch, Hossein Ardeschir A. Ghofrani, Ralf Theo T. Schermuly, Friedrich Grimminger, and Manfred Kaps. Microcirculatory dysfunction in the brain precedes changes in evoked potentials in endotoxin-induced sepsis syndrome in rats. *Cerebrovascular diseases (Basel, Switzerland)*, 23(2-3):140–147, 2007. 121, 123

[165] Karin Rosenkranz and Louis Lemieux. Present and future of simultaneous EEG-fMRI. *Magma (New York, N.Y.)*, 23(5-6):309–316, December 2010. 161

[166] P. Sándor, I. Nyáry, M. Reivich, and A. G. Kovách. Comparative effects of chloralose anesthesia and Sernylan analgesia on cerebral blood flow, CO2 responsiveness, and brain metabolism in the baboon. *Stroke*, 8(4):432–436, July 1977. 173

[167] David J. Sandstrom. Isoflurane depresses glutamate release by reducing neuronal excitability at the Drosophila neuromuscular junction. *The Journal of Physiology*, 558(2):489–502, July 2004. 174

[168] E. Sanna, M. P. Mascia, R. L. Klein, P. J. Whiting, G. Biggio, and R. A. Harris. Actions of the general anesthetic propofol on recombinant human GABAA receptors: influence of receptor subunits. *Journal of Pharmacology and Experimental Therapeutics*, 274(1):353–360, July 1995. 174

[169] Elad Schneidman, William Bialek, and Michael J. Berry. Synergy, Redundancy, and Independence in Population Codes. *The Journal of Neuroscience*, 23(37):11539–11553, December 2003. 125

[170] C. N. Scholfield. Potentiation of inhibition by general anaesthetics in neurones of the olfactory cortex in vitro. 383(3):249–255, 1980. 173

[171] Kristina Schulz, Esther Sydekum, Roland Krueppel, Christoph J. Engelbrecht, Felix Schlegel, Aileen Schroter, Markus Rudin, and Fritjof Helmchen. Simultaneous BOLD fMRI and fiber-optic calcium recording in rat neocortex. *Nature Methods*, 9(6):597–602, May 2012. 5

[172] Adam J. Schwarz, Lino Becerra, Jaymin Upadhyay, Julie Anderson, Richard Baumgartner, Alexander Coimbra, Jeff Evelhoch, Richard Hargreaves, Brigitte Robertson, Smriti Iyengar, Johannes Tauscher, David Bleakman, and David Borsook. A procedural framework for good imaging practice in pharmacological fMRI studies applied to drug development #2: protocol optimization and best practices. *Drug discovery today*, 16(15-16):671–682, August 2011. 162

[173] C. E. Shannon. A mathematical theory of communication. *SIGMOBILE Mob. Comput. Commun. Rev.*, 5(1):3–55, January 2001. 124

[174] Eran Shtoyerman, Amos Arieli, Hamutal Slovin, Ivo Vanzetta, and Amiram Grinvald. Long-Term Optical Imaging and Spectroscopy Reveal Mechanisms Underlying the Intrinsic Signal and Stability of Cortical Maps in V1 of Behaving Monkeys. *The Journal of Neuroscience*, 20(21):8111–8121, November 2000. 176

[175] Robert G. Shulman, Fahmeed Hyder, and Douglas L. Rothman. Cerebral energetics and the glycogen shunt: Neurochemical basis of functional imaging. *Proceedings of the National Academy of Sciences*, 98(11):6417–6422, May 2001. 53

[176] Kenneth Sicard, Qiang Shen, Mathew E. Brevard, Ross Sullivan, Craig F. Ferris, Jean A. King, and Timothy Q. Duong. Regional Cerebral Blood Flow

and BOLD Responses in Conscious and Anesthetized Rats Under Basal and Hypercapnic Conditions: Implications for Functional MRI Studies. *Journal of Cerebral Blood Flow & Metabolism*, 23(4):472–481, April 2003. 174

[177] Yevgeniy B. Sirotin, Elizabeth M. C. Hillman, Clemence Bordier, and Aniruddha Das. Spatiotemporal precision and hemodynamic mechanism of optical point spreads in alert primates. *Proceedings of the National Academy of Sciences*, 106(43):18390–18395, October 2009. 189

[178] Arien J. Smith, Hal Blumenfeld, Kevin L. Behar, Douglas L. Rothman, Robert G. Shulman, and Fahmeed Hyder. Cerebral energetics and spiking frequency: The neurophysiological basis of fMRI. *Proceedings of the National Academy of Sciences*, 99(16):10765–10770, August 2002. 51

[179] S. J. Smith. EEG in the diagnosis, classification, and management of patients with epilepsy. *Journal of neurology, neurosurgery, and psychiatry*, 76 Suppl 2(suppl 2):ii2–ii7, June 2005. 159

[180] Roberto C. Sotero and Nelson J. Trujillo-Barreto. Modelling the role of excitatory and inhibitory neuronal activity in the generation of the BOLD signal. *NeuroImage*, 35(1):149–165, March 2007. 56

[181] Medical I. Specialists. Services. 2013. 158

[182] Danica B. Stanimirovic and Alon Friedman. Pathophysiology of the neurovascular unit: disease cause or consequence?. *Journal of Cerebral Blood Flow & Metabolism*, 32(7):1207–1221, March 2012. 179

[183] H. Stephan, H. Sonntag, H. D. Schenk, and S. Kohlhausen. [Effect of Disoprivan (propofol) on the circulation and oxygen consumption of the brain and CO2 reactivity of brain vessels in the human]. *Der Anaesthesist*, 36(2):60–65, February 1987. 174

[184] Harald Stögbauer, Alexander Kraskov, Sergey Astakhov, and Peter Grassberger. Least-dependent-component analysis based on mutual information. *Physical Review E*, 70(6):066123+, December 2004. 131

[185] Steven M. Stufflebeam. Clinical magnetoencephalography for neurosurgery. *Neurosurgery clinics of North America*, 22(2), April 2011. 160

[186] Keiji Tanaka, Kang Cheng, Hiroshige Takeichi, Tameng Ong, R. Allen Waggoner, Eiji Yoshitome, Shinobu Mizuta, and Kenichi Ueno. Using functional magnetic resonance imaging to study human brain functions, September 1999. 4

[187] Minoru Tomita. Blood flow control in the brain: possible biphasic mechanism of functional hyperemia. *Asian Biomedicine*, 1(1), June 2007. 5

[188] Alessandro Treves and Stefano Panzeri. The Upward Bias in Measures of Information Derived from Limited Data Samples. *Neural Computation*, 7(2):399–407, March 1995. 192, 193

[189] M. Ueki, G. Mies, and K. A. Hossmann. Effect of alpha-chloralose, halothane, pentobarbital and nitrous oxide anesthesia on metabolic coupling in somatosensory cortex of rat. *Acta Anaesthesiologica Scandinavica*, 36(4):318–322, May 1992. 173

[190] M. Ureshi. Stimulus frequency dependence of the linear relationship between local cerebral blood flow and field potential evoked by activation of rat somatosensory cortex. *Neuroscience Research*, 48(2):147–153, February 2004. 6

[191] Hugo Van Aken, William Fitch, David I. Graham, Thomas Brüssel, and Hermann Themann. Cardiovascular and Cerebrovascular Effects of IsofluraneInduced Hypotension in the Baboon. *Anesthesia & Analgesia*, 65(6):565–574, June 1986. 174

[192] Tor D. Wager, Alberto Vazquez, Luis Hernandez, and Douglas C. Noll. Accounting for nonlinear BOLD effects in fMRI: parameter estimates and a model for prediction in rapid event-related studies. *NeuroImage*, 25(1):206–218, March 2005. 75

[193] Kui Wang, Chao Zheng, Chen Wu, Ming Gao, Qiang Liu, Kechun Yang, Kevin Ellsworth, Lin Xu, and Jie Wu. -Chloralose diminishes oscillations in rat hippocampal slices. *Neuroscience Letters*, 441(1):66–71, August 2008. 173

[194] Misfit Wearables. The Great Patent Cliff. 2012. 155

[195] Xingchang Wei, Seung-Schik S. Yoo, Chandlee C. Dickey, Kelly H. Zou, Charles R. Guttmann, and Lawrence P. Panych. Functional MRI of auditory verbal working memory: long-term reproducibility analysis. *NeuroImage*, 21(3):1000–1008, March 2004. 2

[196] Richard G. Wise and Cliff Preston. What is the value of human FMRI in CNS drug development? *Drug discovery today*, 15(21-22):973–980, November 2010. 159, 163

[197] Martin Wolf, Marco Ferrari, and Valentina Quaresima. Progress of near-infrared spectroscopy and topography for brain and muscle clinical applications. *Journal of biomedical optics*, 12(6):062104+, 2007. 160

[198] Mark W. Woolrich, Timothy E. Behrens, and Stephen M. Smith. Constrained linear basis sets for HRF modelling using Variational Bayes. *NeuroImage*, 21(4):1748–1761, April 2004. 76

[199] Mark W. Woolrich, Brian D. Ripley, Michael Brady, and Stephen M. Smith. Temporal Autocorrelation in Univariate Linear Modeling of FMRI Data. *NeuroImage*, 14(6):1370–1386, December 2001. 182

[200] Hong Xuan, Shi-tong Li, and Xin Zhuang. Effects of fentanyl on GABAA

receptor in rat hippocampal pyramidal neurons. *Chinese Journal of Anesthesiology*, 2005. 173

[201] T. Yamakura, L. E. Chavez-Noriega, and R. A. Harris. Subunit-dependent inhibition of human neuronal nicotinic acetylcholine receptors and other ligand-gated ion channels by dissociative anesthetics ketamine and dizocilpine. *Anesthesiology*, 92(4):1144–1153, April 2000. 173

[202] Uma S. Yezhuvath, Kelly Lewis-Amezcua, Rani Varghese, Guanghua Xiao, and Hanzhang Lu. On the assessment of cerebrovascular reactivity using hypercapnia BOLD MRI. *NMR Biomed.*, 22(7):779–786, August 2009. 16

[203] Ying Zheng and John Mayhew. A time-invariant visco-elastic windkessel model relating blood flow and blood volume. *NeuroImage*, 47(4):1371–1380, October 2009. 202